

14-02 | May 9, 2014

# The Application of Visual Analytics to Financial Stability Monitoring

**Mark D. Flood**

Office of Financial Research, U.S.

[Mark.Flood@treasury.gov](mailto:Mark.Flood@treasury.gov)

**Victoria L. Lemieux**

University of British Columbia, Canada

[vlemieux@mail.ubc.ca](mailto:vlemieux@mail.ubc.ca)

**Margaret Varga**

Seetru Ltd., U.K., and University of Oxford, U.K.

[margaret.varga@oncology.ox.ac.uk](mailto:margaret.varga@oncology.ox.ac.uk)

**B.L. William Wong**

Middlesex University, U.K.

[w.wong@mdx.ac.uk](mailto:w.wong@mdx.ac.uk)

The Office of Financial Research (OFR) Working Paper Series allows staff and their co-authors to disseminate preliminary research findings in a format intended to generate discussion and critical comments. Papers in the OFR Working Paper Series are works in progress and subject to revision.

Views and opinions expressed are those of the authors and do not necessarily represent official OFR or Treasury positions or policy. Comments are welcome as are suggestions for improvements, and should be directed to the authors. OFR Working Papers may be quoted without additional permission.

# The Application of Visual Analytics to Financial Stability Monitoring

Mark D. Flood

*Office of Financial Research (OFR), U.S.A.*

Victoria L. Lemieux

*University of British Columbia, Canada*

Margaret Varga

*Seetru, U.K. and University of Oxford, U.K.*

B.L. William Wong

*Middlesex University, U.K.*

May 9, 2014

ALL COMMENTS WELCOME

We gratefully acknowledge helpful comments from Greg Feldberg, Richard Haynes, Roy Kalawsky, Rebecca McCaughrin, Trish Mosser, Mark Paddrik, Peter Sarlin, Del Terry, Wendy Wagner-Smith, and conference and discussion group participants at the August 2011 Records and Information Management for Financial Analysis and Risk Management Workshop at the University of British Columbia, the May 2012 Banff Visual Analytics Interdisciplinary Workshop on Canadian and Global Challenges in Financial Systemic Risk Analysis, the October 2012 IEEE Visual Analytics Science and Technology conference in Seattle, the Open Financial Data Group, a February 2014 research seminar at the Office of Financial Research, and the April 2014 Conference on Data Standards, Information and Financial Stability at Loughborough University. Any remaining errors pertain to the authors alone.

**Views and opinions expressed are those of the authors and do not necessarily represent official positions or policy of the Office of Financial Research, the U.S. Treasury.**

# **The Application of Visual Analytics to Financial Stability Monitoring**

## **Abstract**

This paper provides an overview of visual analytics — the science of analytical reasoning enhanced by interactive visualizations tightly coupled with data analytics software — and discusses its potential benefits in monitoring systemic financial stability. Macroprudential supervisors face a daunting challenge with at least three facets of the financial system. First, the financial system is complex, enormous, highly diverse, and constantly changing. Second, the set of financial and econometric models proposed to help comprehend threats to financial stability is large, diverse, and growing. Third, certain regulatory activities, such as rulemaking and decision-making, generate special requirements for transparency and accountability that can complicate or restrict the choices of tools and approaches. This paper explores these challenges in the context of visual analytics. Visual analytics can increase supervisors' comprehension of the data stream, helping to transform it into actionable knowledge to support decision- and policy-making. The paper concludes with suggestions for a research agenda.

## **Keywords:**

Financial stability, macroprudential supervision, monitoring, systemic risk, visual analytics

# The Application of Visual Analytics to Financial Stability Monitoring

## 1. Introduction

This paper provides an overview of visual analytics and discusses its potential benefits in monitoring systemic financial stability. Macroprudential supervisors face a daunting challenge — the financial system is complex, enormous, highly diverse, and constantly changing.<sup>1</sup> The system generates a seemingly infinite stream of information, arriving at ever-increasing frequencies and spanning a wide spectrum of reliability. The recent crisis demonstrated that stakes are high. Visual analytics has the potential to increase supervisors' comprehension of the data stream, helping transform it into actionable knowledge to support decision- and policy-making.

Managing the risks in financial systems is critical to economic and social well-being. The global financial crisis highlighted the need for enhanced capability to detect, identify, analyze, and understand threats to financial stability. Inadequacies in supervisors' access to and ability to process information hindered an effective response to the crisis. A joint Financial Stability Board (FSB) and International Monetary Fund (IMF) report (FSB-IMF, 2009) notes that, “. . . the recent crisis has reaffirmed an old lesson — good data and good analysis are the lifeblood of effective surveillance and policy responses at both the national and international levels.” Since the crisis, macroprudential supervisors have been working to address data gaps and developing new approaches to financial stability analysis. Beyond the need to create new data sources and analytic approaches, however, the crisis also revealed a need for greater capacity to integrate and make sense of voluminous, dynamic, and heterogeneous financial data.<sup>2</sup>

---

<sup>1</sup> Because we are mostly unconcerned in this paper with the legal details of the supervisory authorities under which formal policies are implemented, we use the following terms interchangeably to avoid cluttering the text with superfluous clarifying language: “financial stability supervisor (or monitor),” “macroprudential supervisor,” and “systemic risk supervisor” refer to national or international authorities responsible for maintaining awareness of and responding to financial-sector stresses and crises with ramifications that extend beyond individual firms or markets. Sarlin (2014) discusses visualization in the context of macroprudential oversight.

<sup>2</sup> For further discussion, see Flood, Mendelowitz, and Nichols (2012), and Flood, Raschid, and Kyle (2010). The present paper originated in a series of interdisciplinary discussions that brought together experts in financial systemic risk and visual analytics. The first of these discussions occurred in May 2012 at the Banff Visual Analytics Interdisciplinary Workshop on Canadian and Global Challenges in Financial Systemic Risk Analysis. A subsequent

One possibility of enhancing the information-processing capabilities of macroprudential supervisors lies in the relatively new field of *visual analytics*. “Visual analytics is the science of analytical reasoning facilitated by interactive visual interfaces.” (Thomas and Cook, 2004, p. 4). The emphasis here is on user interaction, because visual analytics gives human users, with their extensive perceptual and cognitive powers, a central role in a software-assisted analytical process. Keim, et al. (2011) state that “Visual analytics combines automated [data] analysis techniques with interactive visualisations for an effective understanding, reasoning, and decision-making on the basis of very large and complex datasets.” Ultimately, the goal is the “... creation of tools and techniques to enable people to:

- Synthesize information and derive insight from massive, dynamic, ambiguous, and often conflicting data.
- Detect the expected and discover the unexpected.
- Provide timely, defensible, and understandable assessments.
- Communicate these assessment [sic] effectively for action.” (Keim, et al., 2011, p. 7).

A crucial aspect is that the visualizations are designed specifically to support the interactive dynamics (Heer and Shneiderman, 2012) required for users’ analytic involvement with the data in real time.<sup>3</sup>

Visual analytics is part of a spectrum of visualization techniques available to financial stability supervisors, offering varying degrees of user interaction, display animation, and computational analytic capabilities. Visual analytics integrates the most interactive of those tools to support analytical reasoning. A familiar example of visual analytics is the dashboard interactive global positioning system (or GPS) for trucks and cars, in contrast to an old-fashioned paper map. An important difference between visual analytics and other visualization techniques is that visual analytics requires or encourages users to explore data interactively — both the questions and the nature of the answers are unknown. With other visualization techniques, users simply choose what to look at.

---

interdisciplinary panel discussion on Global Challenges in Financial Systemic Risk Analysis: Defining Visual Analytics Solutions, was part of the IEEE Visual Analytics Science and Technology Conference (VisWeek) in Seattle in October 2012. A third group of discussions took place as part of a series of talks hosted by the Open Financial Data Group (OFDG), an informal discussion group concerned with issues of financial data governance and management. The OFDG conversations took place between January 2013 and July 2013.

<sup>3</sup> Contrast this active, real-time user involvement against snapshots of analytical depictions of data, sometimes called “analytical visualizations,” such as the static financial stability maps in Figure 13.

We make three main points in this paper. First, visualization in general and visual analytics in particular are natural aids for financial stability monitors. Supervisors have overlapping mandates, including identifying new sources of financial instability, maintaining situational awareness of developing stresses, implementing decisions and rules that bind the financial sector, and promoting transparency of information to market participants. Visual analytics can support all these objectives, which frequently involve iterative, user-directed search and data analysis. This requirement fits naturally with the user feedback loop of visual analytics (see Figure 10). Typical examples include identifying new sources of financial instability and developing and maintaining situational awareness.<sup>4</sup>

Second, because visual analytics strongly emphasizes comparisons and relationships among empirical data, it is important to define clear, measurable, meaningful abstractions that capture the relevant data semantics and are comparable when applied across the financial system.

Systemic risk analysis is relatively immature in this area and may benefit from work in the cognitive systems engineering community, which has developed techniques for identifying and representing these meaningful abstractions (Woods and Roth, 1988; Rasmussen et al., 2004; Bennett and Flach, 2011). This approach has been applied in a variety of domains, such as air traffic control (Wong, et al., 2007; Vuckovic, et al., 2013) and power grid monitoring (Memisevic, et al., 2005; Sanderson, et al., 2003).

Third, cleaner, better structured data will improve visual analytics. At the same time, visual analytics will also help validate and refine the input data by more quickly revealing what is misleading, missing, contradictory, or not comparable in supervisory data. The development of shared ontologies and data standards can assist this feedback loop. Improving both the visual analytics tools and their input data will therefore typically be an iterative process. We outline a possible research agenda in this direction.

---

<sup>4</sup> “Sensemaking” is the process people use to give meaning to experience. The term sensemaking has primarily marked three distinct but related research areas since the 1970s. In introducing sensemaking to human-computer interaction, Russell, Stefik, Piroli and Card (1993) focus on efficient representations of data to improve understanding and task performance. Dervin (1998) introduces the term to information science as part of a theory of information design. Weick (1995) considers sensemaking in the context of organizational studies. Like Klein, Moon and Hoffman (2006a, 2006b), we focus on sensemaking as the ability or attempt to extract important insights from an ambiguous situation. More exactly, sensemaking is the process of creating situational awareness and understanding in situations of high complexity or uncertainty to make decisions. It is “a motivated, continuous effort to understand connections (which can be among people, places, and events) in order to anticipate their trajectories and act effectively;” Klein, Moon, and Hoffman, (2006a, p. 71).

Applying visual analytics tools to macroprudential supervision is a multi- and interdisciplinary exercise, requiring a clear understanding and definitions of: 1) supervisors' analytic concepts, tasks, and mental models; 2) data; and 3) the visual analytics requirements. Success in developing visual analytics tools is typically achieved through repeated interactions with visual representations of data.<sup>5</sup> Through this iterative process the interactive techniques of visual analytics help evoke novel insights from data.

Visual analytics is only part of the toolkit that macroprudential supervisors will need to transform data into actionable knowledge. Visual analytics does not replace the need for statistical tools, which may be preferable for tasks or phenomena that can be structured with limited dimensionality and that recur frequently enough to provide a reliable sample. Visual analytics seeks to couple interactive visualization tightly with data analysis. Robust statistical models may be less distracted by shiny outliers and can provide formal hypothesis tests and significance bounds. Data mining (Khandani, Kim, and Lo, 2010), also is a promising alternative for high-dimensional or unstructured data.<sup>6</sup> Data mining algorithms are typically highly efficient and can be useful for exploring high-volume or high-dimensional data. The algorithms are typically also designed for generic application and can process unstructured data (Rajaraman, Leskovec, and Ullman, 2014). Because all of these tools continue to evolve, a long-term challenge will be to develop interactive visual interfaces that can support macroprudential oversight in combination with other analytics and reasoning techniques, such as traditional statistics, on-site examinations, network analysis, stress testing, data mining, Monte Carlo simulation, agent-based modeling, etc.

The paper proceeds in three sections. The first provides a broad overview of the use of visualization to depict financial stability or instability. The second section discusses some specific challenges that constrain the use of visualization in this context and provides background for the third part. The third part discusses possibilities for the application of visual

---

<sup>5</sup> In the cheeky formulation of Cherny (2013, p. 16), "Get Yer Stats, Visualize. ... Repeat."

<sup>6</sup> Dimensionality refers to the number of attributes measured for the objects under consideration. Because new attributes can be derived from existing data, it is possible for dimensionality to proliferate. Unstructured data refers to data without the well-defined and consistently applied schemas or constraints on data types, storage formats, and allowable values that make the data amenable to automated analysis. Because all digital data are structured at the lowest (binary) level, "unstructured" refers to the absence of structures layered on at higher levels.

analytics to financial stability monitoring, including a proposed research agenda for the development of visual analytics tools.

## 2. Context: The Use of Visualization in Financial Stability Analysis

Financial stability monitoring, the analysis and synthesis of data and information to identify, understand, and respond to threats to the stability of the financial system, makes extensive use of visual depictions of data and models. Because it covers the full financial system, financial stability monitoring faces a broad range of data and information sources, as well as potential threats and models of those threats. This diversity helps distinguish macroprudential supervision from traditional microprudential supervision and underscores the need for effective tools to navigate through the data deluge.

One way to organize a discussion of financial stability monitoring is to focus on system-level phenomena — liquidity, volatility, concentrated exposures, macroeconomic imbalances, business cycles, etc., — that are common determinants of financial instability. Financial stability maps frequently feature such thematic taxonomies.<sup>7</sup> Another approach, highlighted in Figure 12 and Zhang, (2012, pp. 176-177), is organized around the type and structure of the input data. Because a central element of visual analytics is user interactive exploration of data, we opt to structure the discussion around the broad categories of user tasks, which we group as sensemaking, decision-making, rulemaking, and transparency. We use the other dimensions of systemic phenomena and input data types as a source of illustrative examples.

	Noninteractive	Interactive
Static	No user input after initial rendering, and image does not change. “Fixed.” <i>Example:</i> Newspaper infographic	Ongoing user input, but rendering does not change between input events. <i>Example:</i> Dental X-rays
Dynamic	No user input after initial rendering, but image may change. <i>Example:</i> Animated GIF	Ongoing user input, and rendering may change between input events. <i>Example:</i> Video game

**Table 1: Classification of Visualization Techniques**

<sup>7</sup> For example, the OFR’s Financial Stability Monitor (see Figure 13 and OFR, 2013) is a heat map organized around systemic threats. The IMF’s Global Financial Stability Map (IMF, 2013b, p. 2) is a spider chart organized around a slightly different set of systemic themes.



An extensive literature has analyzed both the craft of visual rendering of data of diverse kinds and the wide-ranging research into the psychology of perception and decision-making based on data visualizations. For example, Tufte (1990, 2001) and Wilkinson (2005) discuss the core principles of data modeling and graphical design for data visualization. Ware (2013) focuses on the psychology of perception as a key factor in the effectiveness of visualization. Lemieux, et al. (2014) and Schwabish (2014) provide a general overview of the application of visualization techniques in economics and finance. Sarlin (2013a) offers a discussion focused on financial stability analytics in particular. Table 1 suggests a simple classification of visualization techniques. “Dynamic” here is synonymous with “animated.” Because static, noninteractive visualizations are so common, we assign a special label, calling them “fixed” renderings.

Interactivity provides an effective means of “playing” with the data by, for example, asking what-if questions naturally while working through a problem. Importantly, visual analytics allows users to redirect the computation of the underlying algorithms by directly manipulating the visual form. For example, a *nonanalytic* interactive visualization might allow a user to add or remove particular data series from a time-series plot. This form of interaction, while dynamic, does not recompute the underlying data. Visual analytics, in contrast, offers the additional ability to act on the algorithms that control the outputs rendered in the visual form. For example, we might instruct the visualization application to combine two securities into a portfolio and calculate or recalculate the risk characteristics of the new entity by dragging one security’s rendering onto the other’s within the time-series plot or in some semantic space to generate and render new results. By performing these transformations in real time, visual analytics extends, develops, and refines techniques for user interaction at the pace of the user’s own cognition.

In this section, we provide a brief overview of the core concepts and some visual approaches to them. Most visualizations of systemic risk are numeric, reflecting the basic fact that financial valuation is a measurement framework. Most financial stability visualizations are also fixed, including the various examples in this section. As we emphasize in our discussion of organizational challenges in Section 3, there are often very good reasons for fixed renderings; however, there are also important cases where interactive visual analytics can contribute significantly.

We focus here on three fundamental aspects of financial stability. The examples touch on some of the enormous range of possibilities for visual renderings and highlight the importance of tailoring the measurement dimensions to both the available data and the concepts being depicted:

- (a) **Concentrated Exposures:** Because concentrated hazards can trigger disruptive business failures, a natural measurement focus is on risk patterns at the level of individual financial firms.
- (b) **Systemic Interconnectedness:** Because patterns in financial instability frequently emerge at the system level via the propagation of credit losses or withdrawal of liquidity support, a natural measurement focus is on the network of financial relationships between key participants.
- (c) **Accumulating Imbalances over Time:** Because large-scale exposures typically build up gradually, a natural measurement focus is to track selected signals over time.

Moving beyond these high-level, fixed renderings to interactive analytics — especially those with variable granularity or “details on demand” — will ultimately require the definition of meaningful hierarchies among concepts or abstractions measured on a comparable scale across the system or some subsystem.<sup>8</sup> We return to this concept in our discussion of the abstraction-decomposition framework in Section 3.

Visualizations, including interactive visual analytics, derive much of their power from by presenting multiple data points on a common measurement scale, facilitating the identification and interpretation of patterns, trends, and anomalies, and encouraging reconciliation across observations. To support these comparisons across the financial system, it is necessary to have stable or “invariant” abstractions — semantically relevant concepts consistently measured across the ecosystem of economic models, institutions, and episodes. Ideally, these abstractions will have both standard machine-readable formats and well-defined semantics. A familiar example of this process is the maintenance of generally accepted accounting principles (GAAP), which provide standard semantics for financial reporting, augmented by the Extensible Business Reporting Language (XBRL), which provides standard formats for automated processing (see Engle, et al., 2013). Similar XBRL-enhanced reporting of bank call reports supports the

---


<sup>8</sup> Consistent hierarchies enable the drill-down interaction (details on demand) of Shneiderman’s (1996, p. 2) mantra: “overview first, zoom and filter, then details-on-demand.”

microprudential supervision of individual firms. The global Legal Entity Identifier (LEI) is another example (OFR, 2013, pp. 82, 97–110).


## Example 1: Concentrated Exposures

A concentrated risk exposure is an unsustainably large contingent obligation or aggregation of obligations that, if triggered, would lead to the failure of a financial firm or system. No consensus definition exists, in part because the scope of exposure matters — firm versus subsystem versus system — and because there are so many often correlated sources of risk and ways to be exposed to them. Visualizations, because of their enormous flexibility in the choice of measurement dimensions and rendering elements, are a natural tool to address this diversity.<sup>9</sup>

### One Important Bank (Koning, 2012)



### Many Banks in Many Countries (IMF, 2013)

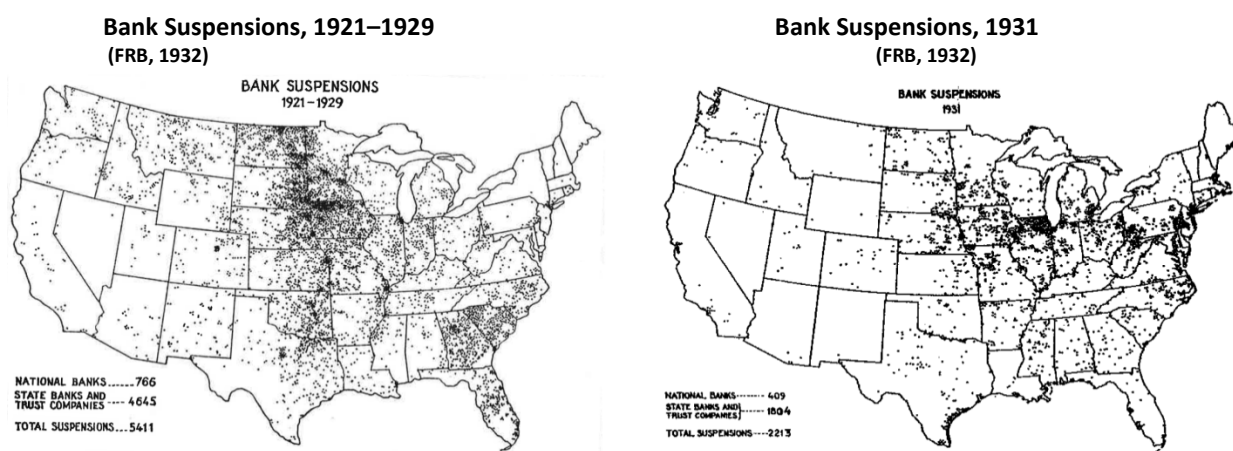


**Figure 1: Risk Accumulation Measured at the Bank Level**

The two panels of Figure 1 highlight intertemporal and cross-sectional comparisons of exposure concentrations. These are noninteractive visualizations. Figure 1 (left panel) depicts balance-sheet information for a single institution, clearly showing China's persistent trade surplus as accumulating foreign currency reserves at the central bank (Koning, 2012). The right panel of

<sup>9</sup> Wilkinson (2005) itemizes the most common rendering elements and describes a coherent implementation framework for combining them flexibly into customized two-dimensional graphics.

Figure 1 is a simple example of how the standardization of accounting data makes meaningful cross-firm aggregations possible (IMF, 2013, p. 17). Risk concentration is conveyed by comparing raw loss exposure (asset quality) to loss-absorption capacity and by the proportion of funding from reliable deposits, as opposed to interbank borrowings. In this case, normalized ratios are aggregated nationally, at a snapshot in time, and compared for a broad cross-section of countries. The three colors, which align precisely with the asset quality metric, clearly reveal a strong correlation with the other three measures. Analogously, the standardization of national income and product accounts was a necessary precondition for tools like the “World Bank Data Visualizer” (see World Bank, 2013b).



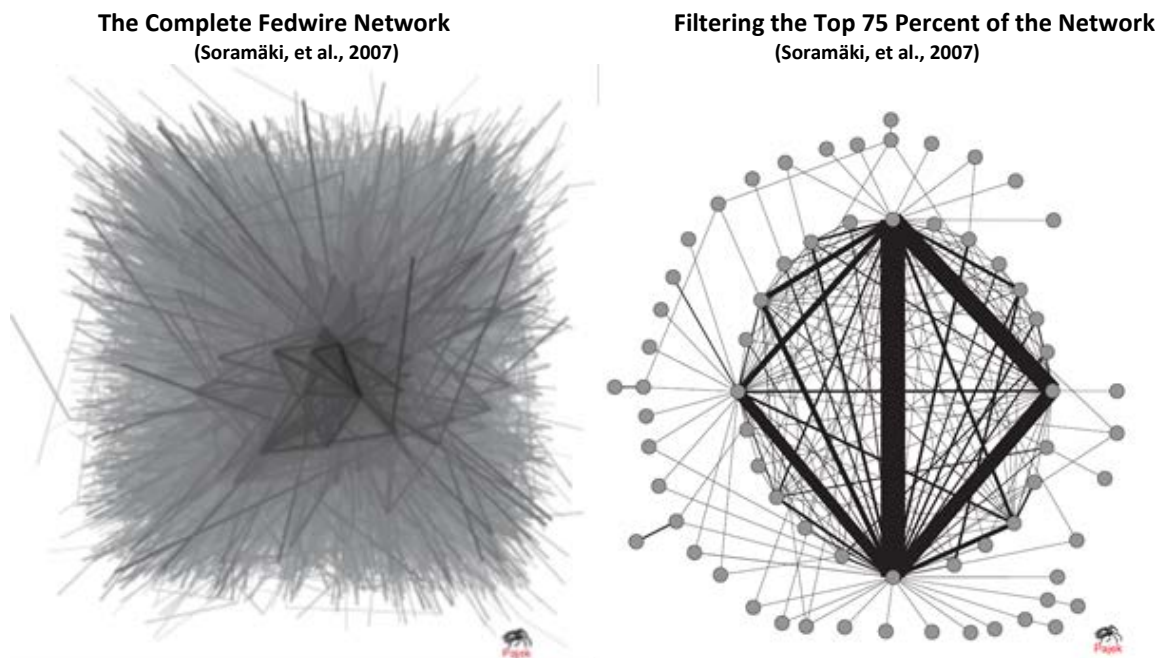
**Figure 2: Geographic Distribution of U.S. Bank Suspensions, 1921–1929 and 1931**

In the domain of financial stability analysis, similar stable abstractions have begun to emerge, although the formalisms are not yet as mature as the GAAP and Call Report standards. For example, as noted, there are no standard, consensus definitions for “concentrated risk exposure.” A step in this direction is the Federal Reserve’s Comprehensive Capital Analysis and Review (CCAR) process, which results in hypothetical bank capital ratios under the stable (across banks) abstraction of a standardized stress scenario. Detailed contract-level data collections from the Y-14 forms support these calculations.<sup>10</sup> At the level of the individual firm, the Basel Committee on Banking Supervision (BCBS) has codified standards for reporting bank capital, risk-weighted

<sup>10</sup> There are different Y-14s at annual, quarterly, and monthly frequencies; see FRB (2012a, 2012b, and 2013) for details about CCAR. In a precursor analysis, the now-defunct Office of Thrift Supervision (OTS) applied a “Net Portfolio Value” model of interest rate risk across the savings and loan industry using standardized inputs and a custom reporting tool. After absorbing the OTS in the wake of the recent crisis, the Office of the Comptroller of the Currency (OCC) abandoned the model as unmaintainable. See OCC, et al., 2011.

assets, and acceptable leverage (BCBS, 2011). Given that stress testing is a counterfactual or “what-if” exercise, it could be a natural candidate for user-driven exploration of the scenario space with a visual analytics tool.

Figure 2 (FRB, 1932, pp. 31–32) compares banking crises in two episodes. Each point on the maps is the location of a bank suspension. The stable abstraction in this case is bank suspension, a binary variable readily comparable across institutions. The geographic pattern effectively conveys the important policy point: The ongoing (in 1932) banking crisis was far more urban and Eastern than the agriculture-related bank failure wave of the 1920s. The inclusion of state boundaries is important as an aid in locating business and financial centers and because banking laws differed significantly from state to state.

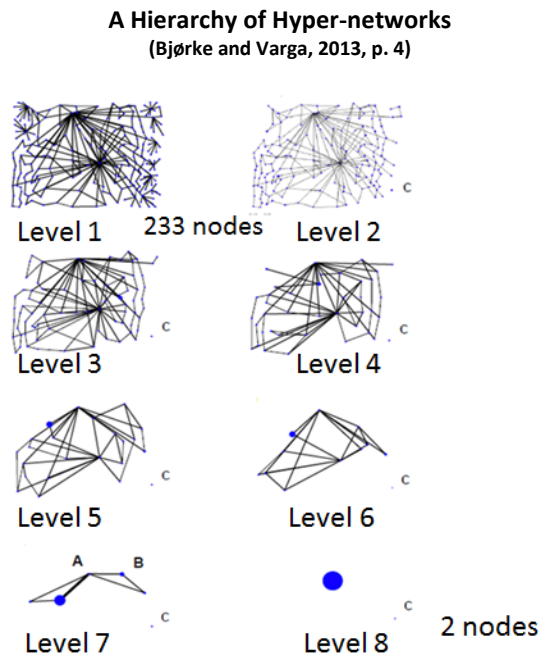


**Figure 3: Topology of the Interbank Payments Network**

### **Example 2: Systemic Interconnectedness**

Stepping back from the microprudential focus of accounting measures, network analysis measures interconnections in the financial system directly. Interconnections matter, because stresses in the financial system can be transmitted quickly from one entity to another through these channels. For example, one firm’s funding difficulties may cause it to withhold liquidity

from others, or joint holdings of the same asset or asset class may create fire-sale exposures if one investor engages in large sales, temporarily forcing down the price for everyone. A natural network analysis application is the counterparty network, with legal entities as nodes and contractual transactions exposures as edges. Figure 3 depicts such a network, in this case for interbank payment flows over Fedwire (Soramäki, et al., 2007). The hairball on the left naïvely shows all participants and flows, while a simple filtering of the graph for the largest nodes generating 75 percent of total payments (right panel) clearly reveals a core-periphery topology typical of dealer markets. The obvious benefit of filtering here suggests a useful dimension for user interaction.



**Figure 4: Network Abstraction – Hypernode**

Alternatively, complex networks can be abstracted by aggregating nodes and links with similar properties into “hypernodes” and “hyperlinks” to generate hierarchies of hypernetworks. This aggregation provides an effective means to reduce visual overload and simplify and enhance the understanding and analysis of complex networks’ interdependencies, robustness, and vulnerabilities (Bjørke, Nilsen and Varga, 2010a, 2010b; Bjørke and Varga, 2013), see Figure 4. The variable levels of abstraction which can be generated using this approach greatly support the user in quickly gaining an understanding of the underlying characteristics of a network; the user


can readily adapt the level of detail / abstraction to the degree desired. An anomaly such as node C, for example, which could be a potential trigger of future events will not be missed using this approach, whereas it could have been removed by simple filtering.

Supervisors have recently begun to receive transaction-level data from swap data repositories, identifying the individual firms participating in trades. For example, Haynes, et al. (2014) use visualizations of confidential transaction-level data on credit default swaps from the Depository Trust & Clearing Corporation to identify trading behavior and patterns in position-taking among dealers in that market. We expect supervisory access to this sort of fine-grained data to expand in the future. The Securities and Exchange Commission's (SEC's) new Market Information Data Analytics System, known as MIDAS, aggregates transaction feeds from a number of exchanges, with node resolution typically at the level of the trading venue (SEC, 2013). The proposed Consolidated Audit Trail would extend node resolution beyond the level of an individual broker-dealer firm handling a trade and further down to the individual customer account or "beneficial owner" (SEC, 2012).

Although generic network algorithms can be useful, tools and visualizations are more powerful when the nodes and edges have domain-specific interpretations. To support visualizations that are adapted to financial data in a systemic context requires special attention to normalization of the input data, which typically come from diverse markets and institutions. As with firm-level accounting reports, standardized network data that capture important economic abstractions are an important building block for higher-level understanding. The LEI, for example, will be invaluable in constructing both visual and nonvisual graphical analyses of financial networks (see Braswell and Mark, 2013; OFR, 2013; Lemieux, et al., 2014; Chan and Milne, 2013). Similar universal identification is needed for the edges of financial networks, particularly counterparty networks. Such identification may come in the form of financial product identifiers. There is a need to address inter- as well as intranetwork analysis (Bjørke, Nilsen and Varga, 2010a) so that we can understand the significance, effect, and correlation of the changes within one network and across other networks.

### Example 3: Accumulating Imbalances over Time

Many aggregate financial stability measures emerge at the level of the system as a whole. Emergent phenomena with systemic implications include liquidity, volatility, concentrated or correlated exposures, macroeconomic imbalances, and business cycles. Economics identified the emergent phenomenon par excellence — the price system — in the late 18th century. Certain key prices, such as the interest rates that compose the yield curve free of default risk, succinctly capture crucial systemic information about the demand for investment and the price of risk.




**Figure 5: Interest Rate Fluctuations in Time Series and Cross Section**

The relationship between long-term and short-term yields reveals much about expected inflation and net returns to financial intermediation. A simple time-series plot of the full curve over time, as in Figure 5 (left panel), conveys a rich history of expansions and recessions, identified primarily by the peaks, which represent expansionary episodes, and troughs, which represent recessionary episodes, in the short-term rate, driven by fluctuations in investment demand (Durden, 2010). The figure also embodies the broad, gradual decline in inflation rates over the period. The near fringe of the surface is the short-term interest rate, providing a succinct synopsis of the Fed’s interest-rate policy over two decades. Other nuances are available only to those with background familiarity with the economic history of the period; this historical context might be made explicit with markers for key events and episodes. As with accounting data described above, the exploration of cross-sectional relationships in interest-rate behavior is similarly well suited to visualization. Figure 5 (right panel) shows changes in interest rates in basis points on



new bank loans between December 2010 and January 2013 (IMF, 2013, p. 10), highlighting the clear divergence between the core and periphery countries in Europe since the financial crisis. The choice of housing and corporate loans rates as the measurement axes was the user's choice. In a visual analytics context, a user might drive these selections via interaction switches.

**Foreign Exchange and Interest Rates,  
Information Dissipation Length, 1998–2011**  
(Quax, et al., 2013)



**Civilian Unemployment Rate, 1979–2013**  
(FRB, 2012)



**Figure 6: Disparate Information Density in Time-Series Plots**

The natural counterpart to the cross-section is the time dimension, which is obviously central to financial calculations. IMF (2011, p. 8), for example, makes time series and cross-section the foundations of its macroprudential policy framework. The venerable time-series plot is ubiquitous in policy briefing books and supervisory reports, and temporal patterns in the data are the crux of the vast, rich field of time-series econometrics and forecasting. Figure 6 reproduces two such plots, with markedly different levels of information density. This disparity reflects the difference, discussed later, between sensemaking exercises, and policymaking and decision-support exercises. In the former, nuance and detail are vital; in the latter, the presence of formal accountability makes lack of ambiguity much more salient. The left panel depicts the information dissipation length (IDL) statistic of Quax, et al. (2013, p. 3). IDL is an abstract, entropy-based measure of the degree of “tight coupling” in a system, with higher values suggestive of instability. The format is presented for a scientific audience and juxtaposes dozens

of time series comprising thousands of daily observations to make a point about the early warning power of IDL. The unemployment rate (right panel) comprises monthly observations on a single key economic indicator over a 34-year interval (FRB, 2012, p. 8). This plot is from the Federal Reserve's *Annual Report to Congress*, and puts the current level of workforce pain in the context of recent business cycles for a rough comparison.

Operational issues affecting temporal data include frequency and timeliness, temporal resolution, and synchronization across data sources. Timestamps from separate sources may not align. Various arrangements have evolved to enable synchronization, such as end-of-day exchange “fixings,” and accounting statements prepared after the fact “as of” the fiscal year-end. These mechanisms are not universal, however. For example, the “Flash Crash” report by the Commodity Futures Trading Commission (CFTC) and SEC highlights problems with network latency in various trading message feeds that may have hampered the price discovery process (CFTC-SEC, 2010, pp. 76-79). More precise timestamps and visual tools for benchmarking message flows from diverse sources might be useful in this context.

### **3. Challenges in Visualizing Financial Stability/Instability**

This section considers the special issues that confront financial stability analysts using visualization techniques. Much of the published work comes from macroprudential supervisors themselves, such as the IMF, Financial Stability Oversight Council (FSOC), Federal Reserve Board (FRB), and central banks in other countries, but vendors, independent analysts, and academics also generate a great deal of published work. By its nature, publication favors noninteractive visualizations. However, as discussed later, this bias is also a function of the relationship between the publisher and the audience. Formal reports, such as the Federal Reserve's *Annual Report to Congress* (FRB, 2012), are not merely exercises in transparency but also part of a larger process of accountability and record keeping. These factors imply a steep penalty on ambiguity and variability that would be unavoidable with animated or interactive visualizations designed for exploratory data analysis and sensemaking.<sup>11</sup>

---

<sup>11</sup> The separate but related challenges of financial records and group decision-making are beyond the scope of this survey. Lemieux (2013) discusses the issues surrounding financial record keeping. Mohammed (2001) discusses the literature on the organizational dynamics of group decision-making, and Heer and Agrawala (2008) discuss

## Challenge 1: The Multifaceted Nature of Systemic Risk

A fundamental challenge is the multifaceted nature of systemic risk. Although there is some consensus about high-level concepts, such as liquidity, leverage, complexity, interconnectedness, etc., the agreement often dissipates with the choice of specific risk measures. In their survey, Bisias, et al. (2012, p. 256) highlight this disagreement:

“One definition of systemic risk is “any set of circumstances that threatens the stability of or public confidence in the financial system” (Billio, Getmansky, Lo, and Pelizzon, 2010). The European Central Bank (ECB) (2010b) defines it as a risk of financial instability “so widespread that it impairs the functioning of a financial system to the point where economic growth and welfare suffer materially”. Others have focused on more specific mechanisms, including imbalances (Caballero, 2009), correlated exposures (Acharya, Pedersen, Philippon, and Richardson, 2010), spillovers to the real economy (Group of Ten, 2001), information disruptions (Mishkin, 2007), feedback behavior (Kapadia, Drehmann, Elliott, and Sterne, 2009), asset bubbles (Rosengren, 2010), contagion (Moussa, 2011), and negative externalities (Financial Stability Board, 2009).”

Others defy the label itself, suggesting that “risk” connotes a reduced-form statistical process devoid of deeper structure. Formulations such as “threats to financial stability” (Berner, 2011), that evoke the underlying economic and institutional mechanisms are preferred.<sup>12</sup>

This ambiguity matters, because visualizations, like any data analysis, must ultimately commit to specific measurement implementations. The challenges of this diversity of economic modeling multiply when the range of analytical approaches meets the many types of financial markets and institutions that affect overall stability. This diversity of definitions is largely a manifestation of the fact that the financial system is itself multifaceted.

The financial sector is a *system* comprising many *components* such as banks and bank holding companies, each with a corporate charter and list of permissible activities that *allocate functionality* to the firm, such as checking services, automatic teller machines, commercial

---

visualization in this context. Savikhin (2013) focuses specifically on the organizational role of visual analytics in the context of decision-making for financial risk management.

<sup>12</sup> The issue of ill-defined problems remains an active area of research. For instance, a recent NATO workshop on visualization for analysis (NATO, 2013) and Roberts, et al. (2014) addressed the topic using cyber security and human terrain analysis as case studies. Roberts et al. (2014) also discussed the issue of provenance in visual analytics at the data, analysis and reasoning levels.

lending, etc.<sup>13</sup> These component institutions, in turn, are collected into subsystems such as “Federal Reserve member banks.” Unlike other familiar systems, such as a minivan or a toaster oven, the financial system is not the result of a single coherent design process, but has evolved as an adaptive hodgepodge of localized innovations and patches as both new opportunities and risk emerge.

We can view the financial system as “complex” in the sense of Sundström and Hollnagel (2011). Complex systems are composed of highly interconnected and interdependent sub-systems. The financial sector too exhibits nested components, often with tight coupling and interesting intercomponent dynamics. In complex systems, global behavior can be difficult to predict, and diverse hazards can emerge from interactions between the various subsystems. The root causes of a malfunctioning system can be difficult to diagnose. In extreme cases, such as the nuclear meltdowns at Chernobyl (e.g. Reason, 1987) or Three Mile Island (e.g. Joyce and Lapinsky, 1983), total system failure can result. Given the potential for surprise events, nuanced situations, intricate interactions, and unintended consequences of system design choices, it is not unusual to have human experts either in the loop (e.g., ballistic missile launch silos) or at the ready (e.g., airliner cockpits in autopilot mode) in critical complex systems. In the context of systemic financial events, for example, it is commonplace to involve senior policymakers who have both expert professional judgment as well as the legal authority to intervene.

Cognitive systems engineering is a formal approach to including human experts into engineered systems, focusing on the functional relationships and constraints among the components of the system—for example, what must not break to avoid disaster (Rasmussen et al., 1994; Vicente, 1999; Woods and Hollnagel, 2006).<sup>14</sup> Ecological interface design (Vicente and Rasmussen, 1992; Burns and Hajdukiewicz, 2004) is a related subfield devoted to human user interfaces that

---

<sup>13</sup> The italicized terminology here comes from systems engineering. A well designed system allocates to specific components clear “ownership” for particular bits of functionality. See Blanchard and Fabrycky (2010) for a textbook treatment.

<sup>14</sup> A “faceted” approach is another way to representing the relationship between system components. Experts identify key subsets of component or system attributes to be tracked together as facets. The faceted approach originated in the social sciences (Guttman, 1959), where it has been applied to the different facets of relationships within a community (Ginges and Cairns, 2000), including interactive visualizations (Rodrigues, et al, 2011). In a visualization context, facets are mapped to particular rendering elements, such as coordinate axes, glyphs, colors, etc. A systematic faceted analysis can help distinguish between the risk of failure of the components and subsystems, the possible causes of failure, such as the role of humans in the system, and the consequences of a failure.

represent functional relationships and constraints “ecologically”—i.e., realistically and holistically—so the user can understand how the system as a whole is functioning, rather than the components and subsystems in isolation. A systems perspective can help in the design of monitoring and control tooling to track the performance and interaction of specific components or subsystems to enhance sensemaking, situational awareness or decision-making, as the task requires. In a macroprudential context, a cognitive systems design could provide supervisors with a functional overview across the components of the system, together with the ability to drill into an individual component to understand its detailed status and how localized stresses might expand to propagate through the system.

The challenges that a diversity of economic modeling presents tend to multiply when the range of analytical approaches meets the various markets and institutions that compose the financial system. The problem is further compounded by the fact that the system changes constantly as participants innovate and risks evolve. As a result, no one analytical solution, including visual analytics, will address all requirements. Human judgment will remain an integral part of the analytical process. Computational tools assist in these efforts by performing certain data manipulations much more rapidly than human experts can, making an enormous task more manageable. In this context, visualizations can provide a powerful interface for connecting human users as “components” in an analytical system or process.

## **Challenge 2: Acquiring the Right Types and Adequate Quality of Data**

Sometimes the display of financial conditions can mislead users due to data quality problems afflicting the underlying sources. Accounting data have built-in processes for data scrubbing, such as double-entry bookkeeping, audits, and supervisory examinations. Exchanges typically publish price data only after an extensive trade confirmation and settlement process. Financial stability analysts cannot assume that all their data sources will be so pristine. For example, data quality and underwriting standards for subprime mortgages were significant factors in the recent crisis (Keys, et al., 2010; Hunt, et al., 2013). Financial stability analysis may require access to long time series. In many cases, data may be missing values for certain periods, the semantics of the values will have changed, or human data entry errors or operating problems in systems infrastructure will have affected the reliability of the data. Supervisors will lack key data in

certain cases. Indeed, the phrase “shadow banking” is a nod to the large risk exposures lurking outside of supervisory scrutiny. Participants and regulators alike are hampered by opaque assets, complex networks of securitization, and rehypothecation of collateral, model risk, and abrupt regime shifts in the behavior of prices and their correlations. As trading feeds begin to deliver to supervisors the high-frequency message traffic of quotes, orders, corrections, cancellations, and unconfirmed transactions, new data quality challenges will appear (Dick-Nielsen, 2009). The Group of 20 (G-20) has initiated a process to identify data gaps (FSB-IMF, 2009 and 2012), and the Bank for International Settlements will be hosting a “data hub” to improve supervisors’ visibility across national boundaries (Caruana, 2012). Visual analytics tools will need to provide the reader with clear cues regarding the accuracy, reliability, uncertainty, and availability of the data inputs — or lack thereof.

In addition, there are basic operational challenges in managing diverse data sources from across the financial system. The scope of coverage for a particular economic or financial variable is an important dimension of the problem. Individual risk exposures may be neutralized through hedging or diversification in a larger holding company or sovereign aggregate — or they may be compounded by similar exposures elsewhere. Managerial accountability for firms, political jurisdictions, and currency zones are also closely tied to the scope of authority. The ability to identify managerial entities using LEIs and connect them accurately to risk exposures is crucial. In practice, the granularity of measurement for the latter is variable, with some exposures and other attributes measured at the position level, and others only at the portfolio, branch, firm, or national level. A common artifice is to aggregate exposures into buckets, to reduce memory footprints, and smooth gaps and outliers in the data. Choice of aggregation technique, however, may have a significant effect on how an analyst perceives risk exposures, so caution is warranted (Lemieux et al, 2014).

### **Challenge 3: Visually Representing Underlying Concepts and Processes**

A key challenge in using visualizations to communicate about financial stability lies in designing good representations for core concepts and their relationships. Because visualization relies on comparison, good design requires the identification and refinement of stable abstractions — values, typically numeric, that reflect or illuminate the relevant concepts while remaining

commensurate over time and across entities in the system. Good design also requires the selection of good visual representations of the measured entities, concepts, and relationships. Visual analytics aims to go beyond fixed renderings to link visual representations dynamically with algorithms so the analyst can use computation to steer the algorithms in a deliberate manner. For example, in subspace clustering, a data mining technique used for multidimension reduction, an analyst is able to steer the clustering algorithms to reveal low frequency but potentially interesting dimensions that would otherwise be overlooked by automatic data mining procedures (Tatu et al., 2012). Not only must visual analytics systems render well the core concepts and processes, they must also have carefully calibrated interaction techniques that support manipulating the underlying algorithms and the associated data spaces through interaction with graphical elements.

Some key steps in the design of visual analytics systems are:

1. Determining what to represent.
2. Choosing visual forms to represent objects.
3. Designing underlying computational algorithms.
4. Choosing the interactions that connect the visual representation with the underlying analytics.

The way these steps are executed will significantly influence the usability of the technology. In this section, we will briefly discuss the issue of what it is that we should represent, which will guide us in determining what to design.

An example of a common framework for decomposing complex systems for visual analytics implementations is the *abstraction-decomposition* space (Rasmussen et al., 1994; Vicente, 1999). To apply this framework, a visualization designer would deconstruct financial systemic risk analysis by identifying the different levels of abstraction along two dimensions.<sup>15</sup> The advantage of representing the financial system within such a formal framework is that it allows financial

---

<sup>15</sup> Because the financial system is extremely complex, for the purposes of this paper we provide only a very cursory discussion of how the system might be decomposed to show where potential systemic risks might emerge. This overview is only to illustrate generally how the abstraction-decomposition space might highlight the potential benefits of visual representations for macroprudential analysts. No claim is made that the decomposition analysis here is rigorous – that is a topic for a future stream of research.

stability experts to describe domain-specific concepts in a way that is accessible to visualization designers and software engineers.

The abstraction-decomposition space comprises two orthogonal dimensions: the *decomposition hierarchy* and the *abstraction hierarchy*:

1. The decomposition hierarchy consists of three levels of decomposition:
  - a. The whole system
  - b. Subsystems
  - c. Components
2. The abstraction hierarchy consists of five levels:
  - a. functional purpose describing what the system is intended to achieve in the work domain;
  - b. abstract function describing important system concepts, such as liquidity or value of financial products;
  - c. generalized function representing the subsystems and their component parts (for example, indicating how they provide the necessary outputs to calculate liquidity and value);
  - d. physical function that represents the components of the system and their states (for example markets, counterparties, intermediaries, etc.); and
  - e. physical form referring to the configuration, and location of components and subsystems.

Achon and Jamieson (2003) provide an example of an abstraction-decomposition hierarchy for a portfolio management system, depicted in the left panel of Figure 7. They magnify and decompose only one branch of portfolio management, immediately making apparent that many different visual forms are needed to depict the performance of financial portfolios. Figure 7 also shows components (“Enablers” in the diagram) functionally related to “Processes” that perform the “Income transfer” function. In a financial stability context, one example for an abstraction-decomposition hierarchy might be the nested ownership of subsidiaries in financial holding companies. Each holding company would be a subsystem containing a tree with individual subsidiaries, branch offices, or business units as the leaf nodes. An obvious set of functionality to



model in the abstraction hierarchy would be the permissible activities granted to each subsidiary under its corporate charter. However, although straightforward, simple lists of legally permissible activities will typically not correspond directly to crucial financial stability factors, such as credit exposure, leverage, or liquidity.

Once a problem domain has been structured in this way, techniques of human-computer interaction design can represent the relevant functional relationships visually by mapping key component attributes, facets, derived risk measures, etc., to particular rendering elements.<sup>16</sup> The goal is to present the “terrain” of the system, where elements in the visual display are carefully juxtaposed to highlight meaningful comparisons and relationships between the various processes, components, and derived metrics. Such a presentation requires a sympathetic understanding of the subject matter domain. For example, Vuckovic, et al, (2013) untraditionally represent a key functional relationship. They show potential air traffic conflicts as relative position vectors that make the separations between pairs of aircraft visible, instead of predicting likely collisions based on flight plan data.

Another example of a structural modeling hierarchy in finance is the “unified model” of Brammertz et al. (2009), which decomposes financial contracts into bundles of promised cash flows that follow a relatively small number of standardized patterns. These cash flow patterns interact with an event stream over time to generate realizations of the risky exposures. Brammertz (2013) argues the primary source of complexity in the system arises from the financial contracts that connect financial entities. As depicted in the right panel of Figure 7, financial contracts, buffeted by events along the timeline, yield a value and liquidity that can be expressed in terms of risk. The unified model provides an abstract description of the components and subsystems that make up the financial system, projecting contracts, market behaviors,

---

<sup>16</sup> For a good textbook treatment of human-computer interface design, see Shneiderman, et al. (2009). One particular design paradigm, ecological interface design (Burns and Hajdukiewicz, 2004), is well suited to ongoing financial market monitoring, where repeated tasks are commonplace (hourly price updates on key markets), but user attention must contend with manifold distractions, such as news alerts, office banter, etc., while the representation design paradigm (Bennett and Flach, 2011) provides human factors and ergonomics guidance on how to map semantically relevant process variables to visual renderings.

income, revenues, and liquidity into abstract cash-flow patterns and relationships for a composable representation of financial risk.<sup>17</sup>

An abstraction-decomposition hierarchy provides clear guidance on how to structure a problem domain to make it accessible to visual analytics. At the same time, a moment's reflection reveals the daunting nature of this challenge. None of the elements of either hierarchy is defined consistently across the financial system. For example, even assuming that a legal entity is the finest-grained component resolution needed, implementing LEI is in the infancy stage. Central concepts, such as "liquidity" and "risk" have literally scores of concrete definitions, with little consensus on which to use in what context. Despite that problem, the organizing principles of the abstraction-decomposition hierarchy provide good guidance on the structural information needed for a coherent, structured description of a financial system or subsystem. We can use this framework to help outline a research agenda around data and modeling gaps for the financial system.

#### **Challenge 4: Organizational Issues**

The foregoing examples hint at the forces affecting the use of visualization for macroprudential analysis. Macroprudential supervisors convert data on institutional and financial market conditions into analyses, such as briefings, reports to Congress, research working papers, and staff studies, etc., and formal decisions, such as enforcement actions, Federal Reserve open-market operations, prompt corrective action interventions, etc.. Broadly, macroprudential tasks tend toward one of two poles: 1) general monitoring by which authorities track potentially stressful conditions in the financial sector, and 2) formal supervision, which implies the possibility of regulatory enforcement actions. The former focuses on situational awareness, where tools and techniques should be as diverse and flexible as possible so analysts understand evolving conditions, but the latter focuses on formal decisions with tangible consequences, where tools and techniques should be streamlined for clarity and after-the-fact accountability. As

---

<sup>17</sup> As a nonfinancial analogy, a nuclear power plant would typically have a high-level function model of heat and mass exchanges that are abstracted from the hot water and cold water reservoirs; see Memisevic et al. (2005), and Sanderson, et al. (2003).

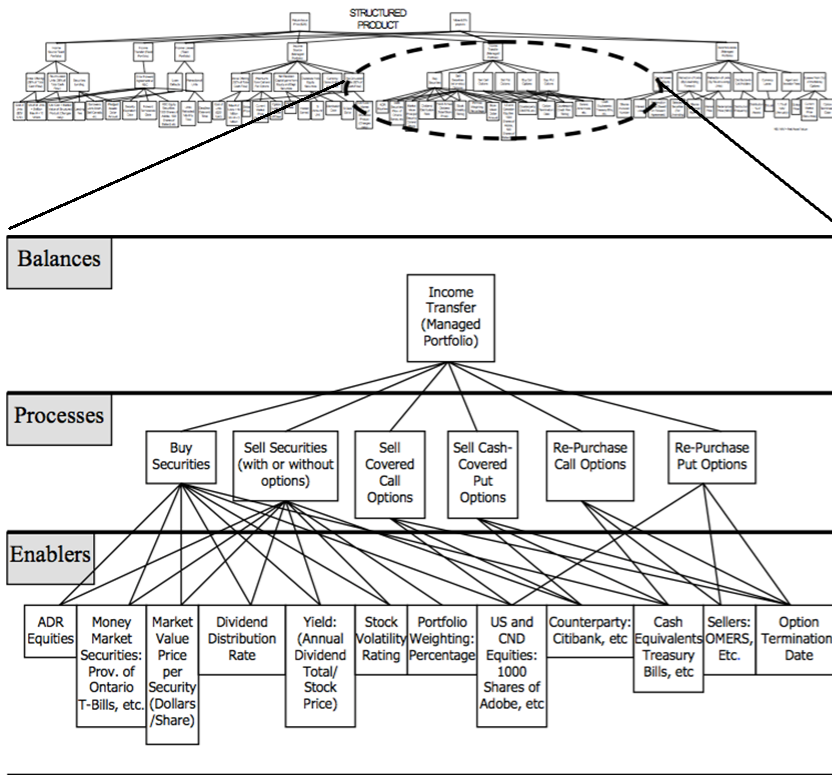
a simple organizing framework, we suggest a high-level breakdown of core supervisory functions:<sup>18</sup>

- Sensemaking
- Decision-making
- Rulemaking
- Transparency

---

<sup>18</sup> Note that we are redefining the terms “sensemaking” and “decision-making” here relative to their traditional use.

### Abstraction Hierarchy for Portfolio Management (Achonu and Jamieson, 2003)



### Base Architecture for Mapping Risk Factors to Risk Measures (Brammertz, 2013)

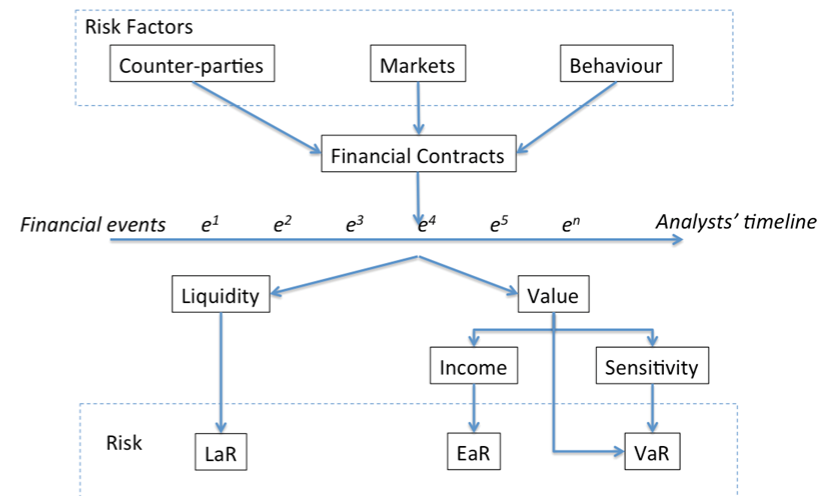
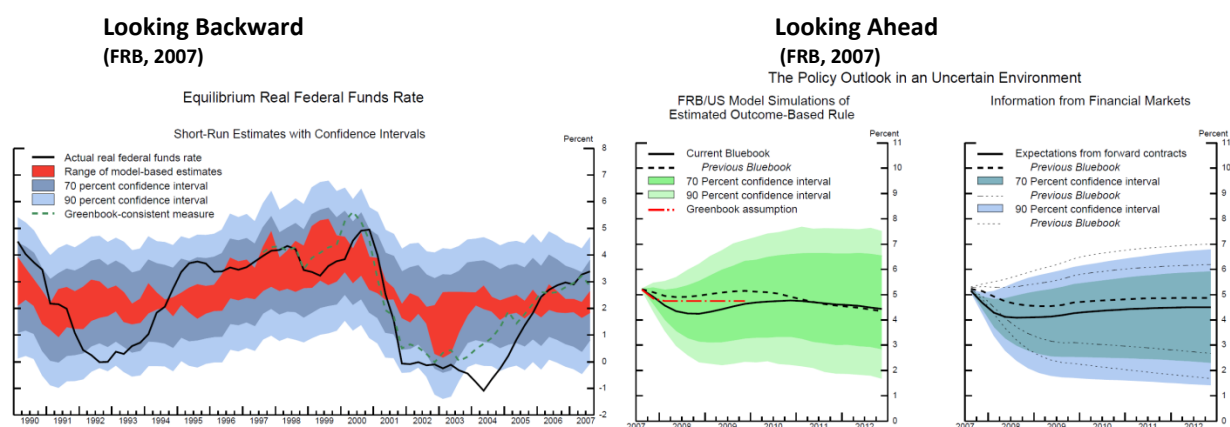


Figure 7: Two Examples of Abstraction Hierarchies in Finance

Macroprudential monitoring inevitably involves a sensemaking exploration of “uncharted territory.” The financial system evolves and innovates to exploit real or imagined new profit opportunities, often creating real or imagined new risk exposures in the process. Supervisors must maintain significant research capacity for this purpose. Sensemaking for new or unusual circumstances is the detailed front end of supervisors’ information-management challenge, with subject-matter experts, such as data specialists, economists, and bank examiners, focused on exploratory analysis of specific events or issues to gain a nuanced understanding of conditions. Data are fine-grained and often raw. Analytical tasks are tactical and often opportunistic. Trial-and-error analyses may be prevalent, and deadlines may be tight, precluding refined investigations. Sensemaking capacity is especially crucial for macroprudential monitors because of the possibility of financial crises. A financial crisis episode is the supervisory “perfect storm”: the stakes are high, time frames for decision and action are likely to be compressed, and familiar statistical and economic relationships in the data are unreliable or inadequate. The interactive techniques of visual analytics are likely to be especially fruitful at this level. In later sections, we focus on the sensemaking challenges for financial stability monitoring.

Decision-making is an operational role. Decisions are rendered within the bounds of existing authorities and rules. For example, the Federal Open Market Committee meets every five to six weeks to implement the Federal Reserve’s monetary policy (FRB, 2013b) by deciding on open market operations. Formal decision processes typically generate agendas and minutes for the public record. Decision-makers often receive formal or informal advance briefings; the “arbitrary and capricious” standard for accountability (Watts, 2009) implies a need for solid analysis and strong documentation. One benefit is the reduction in uncertainty from a conversion of complex, subjective, and ambiguous information into a clear ruling. To support accountability and aid recordkeeping, fixed visualizations are generally preferred at this level because interactive visualizations are difficult to capture and preserve as evidence of inputs to decision-making. Such systems can be designed to alert or guide an analyst to avoid biases or problematic tendencies (Kahneman, 2011; Hutchins, 1995). Such designs, however, are predicated on a careful analysis of decision-makers as they make decisions in real life situations, using such techniques as cognitive task analysis, verbal protocol analysis, and pair analysis (Trickett, et al. 2007; Crandall, et al. 2006; Schraagan, et al. 2000; Arias-Hernandez, et al, 2011).

Rulemaking is a strategic role, in which formal authorities of supervisory action are defined or refined. The primary examples are legislation and regulation. The process of introducing or modifying regulations is highly formalized and open to public scrutiny, which is often extensive. In the United States, notices of proposed rulemaking are published in the *Federal Register* (OCC, et al., 2006), followed by extended periods of public comment; iterations of proposed rules may be repeated as appropriate and may be preceded by an advance notice of proposed rulemaking. Because the law does not have formal structures for interpreting and applying nontextual rules, visualizations tend to be rare at this level, even as supporting analysis. When used, visualizations are most often fixed. There is a potential to apply text visualization to the analysis and understanding of textual documents (Chuang, Manning, and Heer, 2012).

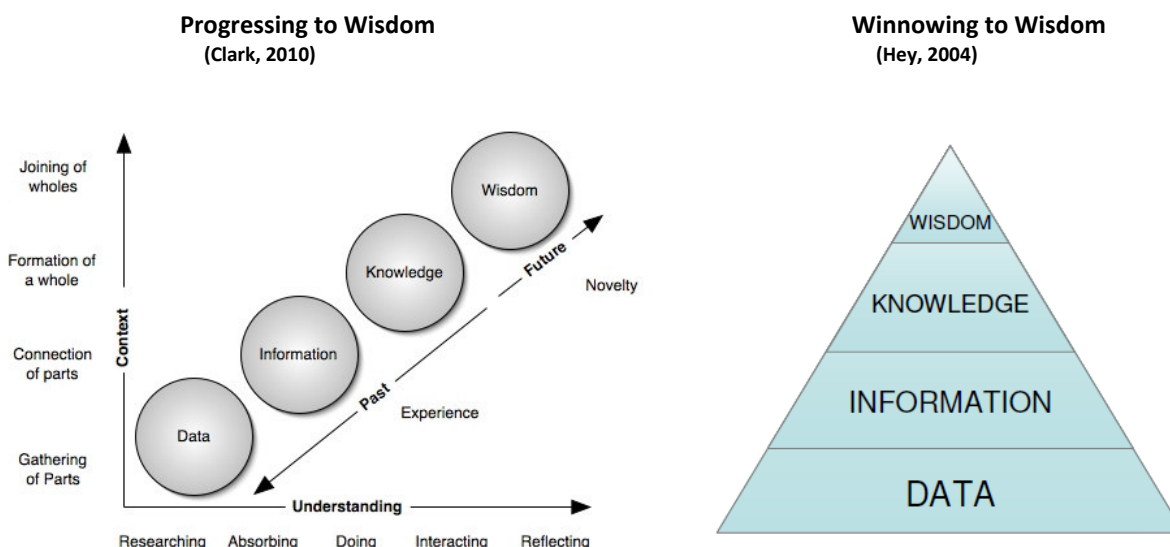


**Figure 8: Model Uncertainty and Decision Support**

Macroprudential agencies pursue transparency through the publication of formal reports, technical analyses, and economic and financial statistics. Transparency can reduce asymmetric information and policy uncertainty, facilitating coordination and contracting among market participants.<sup>19</sup> Publications of formal reports and policy deliberations are important inputs into the accountability process for supervisory authorities. For example, Figure 8 reproduces two depictions of the Federal Reserve’s confidential briefing materials (FRB, 2007, 17 and 21) for the Federal Open Market Committee as the recent financial crisis was beginning in September

<sup>19</sup> For a recent survey of transparency issues in a supervisory context, see Flood, et al. (2013). For a discussion of the costs of policy uncertainty (and benefits of reductions), see (for example) Amengual and Xiu (2013) and Pástor and Veronesi (2013).

2007.<sup>20</sup> The left panel compares the realized historical rate (black line) with the dispersion of (red range) and confidence intervals around (blue ranges) the contemporaneous model estimates. The right panel shows the best-estimate forecast (black lines) and confidence intervals (blue and green ranges) for the federal funds rate in the future. The figures do a good job of presenting the uncertainty that surrounds many policy decisions.<sup>21</sup> Like the Federal Open Market Committee briefings, other transparency publications include visualizations, typically tailored to a particular constituency.<sup>22</sup> Briefing materials like the Federal Reserve’s Bluebook (FRB, 2007) provide the decision-making committee with a core set of common knowledge. It is important that these materials employ *fixed* visualizations, because this provides a basis for common knowledge and clarity — all committee members see the identical image, and staff can explain precisely how it was constructed — for both ex-ante decision support and after-the-fact accountability.



**Figure 9: Two Perspectives on Data-Information-Knowledge-Wisdom**

<sup>20</sup> The figures come from the Federal Reserve staff’s Bluebook of monetary policy alternatives. The confidential Bluebook is made public after five years to balance the pressures for transparency and confidentiality in policymaking.

<sup>21</sup> In fact, as the crisis unfolded, the Federal Reserve dropped the funds rate to essentially zero in late 2008, where it has remained since. Hindsight is 20-20.

<sup>22</sup> Especially useful for financial stability analysis are the annual reports of the Office of Financial Research (OFR, 2013), the Financial Stability Oversight Council (FSOC, 2013), the Federal Reserve (Fed, 2013), and the semiannual financial stability reports of the Bank of England (BoE, 2013) and the IMF (IMF, 2013). Noteworthy from a visualization perspective are the online tools providing access to downloadable data and visualizations with varying degrees of interactivity, such as the Federal Reserve Economic Data site (FRED, Fed-St. Louis, 2013), the IMF Data Mapper® (IMF, 2013b), or the World Bank’s DataBank (World Bank, 2013).

Both the decision-making and rule-making roles have a preference for fixed visualizations because of the need for evidence of decision-making for accountability. The requirement for fixed renderings can prevent decision-makers from exploring options and simulating decision outcomes, however. To overcome this limitation without sacrificing accountability, dynamic visual analytics systems must have the capacity to generate trustworthy evidence (Lemieux and Dang, 2013). An alternative is techniques for tracing decisions through a visual analytic process, known as *analytic provenance*. More commonly encountered in the world of archives, art and antiques, provenance is the process for tracing changes to an entity to compare its original state with its present state. This can enable tests for authenticity (i.e., is the entity what it purports to be) and integrity (i.e., has it been altered inadvertently or deliberately from its original state).

Provenance is also about identifying and preserving the antecedents and context of an object, such as a decision, so that it can be properly understood and evaluated. While it is currently possible, though not unproblematic, to trace changes to data (data provenance) and the processes applied to effect those changes (process provenance), it is much more difficult to trace the human analytical reasoning process. A number of visual analytics researchers are working on this problem. Nguyen et al. (2014a, 2014b) are investigating a set of techniques known as SchemaLine and TimeSets that enable a user to construct streams of explanatory narratives from information automatically extracted by the system, while providing annotation capabilities to trace the choices made in the evolving story. Other approaches include techniques for tracking the analytical processes, the order in which they occurred, and annotating changes in one's analytical considerations (Gotz and Zhou, 2008; Kadivar et al 2009). These stages in the process can be replayed or rearranged. When replayed, the system repeats the full computation, even allowing the analyst to change the values and parameters (Walker et al, 2013).

In practice, of course, authorities engage in a spectrum of activities that mix various monitoring and transparency tasks with formal interventions and new rulemakings. This transformation of data into knowledge and action suggests the familiar data-information-knowledge-wisdom (DIKW) hierarchy, depicted in Figure 9. The four levels of the hierarchy are complex concepts, and there is vigorous debate around both terminology and meaning of each level.<sup>23</sup> We focus on

---

<sup>23</sup> For a textbook exposition of Data-Information-Knowledge-Wisdom, see Sarkar (2013). For examples of the terminological debate, see Hey (2004), Rowley (2007), and Frické (2009).



one aspect of the hierarchy, applied narrowly to the series of transformations that convert raw input data into a final two-dimensional image rendered to pixels on a display device or printed page.

The key distinction is the extent to which information is lost while stepping through the hierarchy. Simply, in the left panel of Figure 9, the balls retain the same size at each stage; in the right panel, each stage discards extraneous information to converge on the core truths. The sharp dichotomy in information density between the two panels of Figure 6 exemplifies the issue. The right panel in Figure 6 distills the input data down to a single key abstraction, civilian unemployment, normalized to adhere to a stable scale and plotted monthly over time; everything else is omitted as superfluous. The left panel presents two versions of the information dissipation length for interest rate swaps and exchange rates over time along with an enormous amount of contextual information from both markets. Clearly, the authors of IDL believe the additional context will be useful for their readers; the authors of the unemployment graph do not. This choice is reasonable in both cases: IDL is a complicated calculation, newfangled and abstract. Unemployment is a simple aggregation, familiar from the nightly news and quite concrete; most readers will have personally experienced an episode of joblessness at some point in their lives.<sup>24</sup>

The distinction is important for financial stability monitoring, because there is no consensus yet on a canonical set of familiar abstractions that are the “correct” way to measure systemic fragility. Bisias, et al. (2012, p. 256) emphasize that a “robust framework for monitoring and managing financial stability must incorporate both a diversity of perspectives and a continuous process for re-evaluating the evolving structure of the financial system and adapting systemic risk measures to these changes.” There will always be new emergent risks and approaches to their identification. In other words, financial stability monitors do not have the luxury of optimizing fixed tasks in a relatively stable operating environment, like air traffic control (Wong, et al., 2007) or professional sports (Pileggi, et al., 2012). Financial systemic risk analysis tasks are not performed by a single individual working alone on deterministic tasks. Instead, risk analysis involves the collaboration of many individuals across organizational and, increasingly,

---

<sup>24</sup> A key constraint driving these visualizations to opposite extremes is that they are both rendered for a static print medium. Given that there must be a single rendering, one has maximized clarity, the other has maximized context. Neither is a necessarily bad choice, given their respective audiences. An interactive medium could finesse the dichotomy by highlighting only the most important series while offering details on demand.

geopolitical boundaries working together under often uncertain conditions. The environment in which financial systemic risk analysis is carried out is highly complex. These constraints require that solutions to enhancing financial systemic risk analytic capabilities be flexible and able to deal with unforeseen circumstances, variable tasks, and integration, presentation and interpretation of large amounts of uncertain, and incomplete and contradictory information that degrades over time.

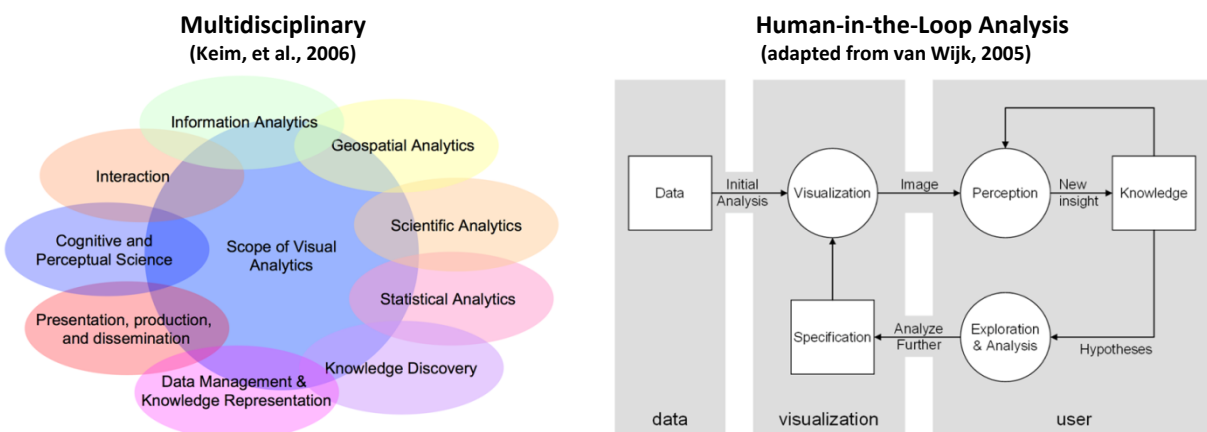
Many of the tasks in financial stability analysis that would benefit most from visualization techniques are exploratory sensemaking tasks, often performed ad hoc, with a human in the loop. Visual analytics, which ideally blends the best of human and computational approaches, has the potential to address the information challenges in this domain by combining humans' judgment and pattern-recognition strengths with the machine's raw calculating power. Though computation can help address the information processing requirements, for example, for text mining of contract data or machine learning for aggregation of systemic risk indicators, it is insufficient on its own. Computational approaches require certain preconditions: that machines can compute the optimum clustering of data; that data are certain, clearly structured, and semantically well defined; and that data are valid, complete, relevant, correct, up-to-date, and do not change over time. It also requires that domain problems are well defined and clearly specified. In financial systemic risk analysis, these preconditions are rarely met. Clustering algorithms still require human assistance to determine the optimum clusters. Data structures and semantics are often ill defined and ambiguous, requiring a human to interpret meaning. Data are incomplete, corrupt, incorrect, contradictory, out-of-date, uncertain, and change over time, again requiring human intervention to interpret, clean, or preprocess them.

## **4. Visual Analytics as a Tool for Financial Stability Monitoring**

### **Overview**

Visual analytics tools and techniques can provide effective means for extracting information and deriving insight from massive, dynamic, often ambiguous and, at times, contradictory data. Visual analytics' strengths lie in not only detecting the expected but more importantly in discovering the unexpected (Thomas and Cook, 2005). Visual analytics can provide timely,

defensible, and clear understanding and assessment of data and situation and communicate that assessment for effective decision-making. Visual analytics draws upon many different areas including, for example, analytical reasoning techniques, visual representation and interaction techniques, data representation and transformation, and techniques to support production, presentation, and dissemination of the results (see left panel in Figure 10). The high-level goal of visual analytics is to combine the visual and cognitive intelligence of human analysts, such as pattern recognition or semantic interpretation, with machine intelligence, such as data transformation or rendering, to perform analytic tasks iteratively. In visual analytics, this feedback loop (see Figure 10, right panel) operates through interactive visual interfaces.



**Figure 10: How Visual Analytics Works**

Visual analytics has certain advantages over traditional statistical methods of transforming large volumes of heterogeneous data into actionable knowledge, because humans have evolved exceptional visual and spatial skills that include the ability to detect edges and discontinuities, exceptions and outliers, and patterns. Visual and spatial attributes, such as color, shape, and motion, can be transformed into a graphical image to provide a rich visual description of data. When features can be perceived before conscious attention (“pre-attentively”), they are understandable at a glance, and much more rapidly than words.

By encoding data and functional relationships (see the previous section on the abstraction-decomposition hierarchy) into images that stand out (Treisman, 1985) in the human field of vision, visualizations help shift cognition to the perceptual system, with visuals acting essentially as a form of externalized memory. Such externalized representations can enlarge an analyst’s

problem-solving capabilities by enabling the processing of more data without overload. As previously mentioned, effective visualization can be used to help overcome cognitive biases that can prevent effective risk-based reasoning. Visual cues can help analysts understand where biases arise because graphical representations stand out forcefully in human perception. Visual analytics exploits these general strengths of visualization by using a visual interface to connect human experts as “components” in a larger analytical system.

There has been terminological debate over the respective definitions of “visual analytics” and “information visualization” (for example, Yi, et al., 2007), centering on the scope and nature of the analytical models available to the human user in the interaction loop. At one extreme lie fixed visual displays, such as the examples in Figure 8. As we have noted, noninteractive visualizations play a necessary role in the supervisory process, especially for the most important decisions and reports.<sup>25</sup> At the other extreme, sophisticated algorithms and analytics are available for the user to interact directly and dynamically with the data and underlying algorithms through the visual form (for example, in the bottom interaction pathway in the right-hand panel of Figure 10) (Heer and Shneiderman, 2012). Some have referred to this form of interaction as “direct data manipulation” (Wang, et al., 2012). Some prototypes provide additional capabilities. One, called RiskVA (Wang, et al., 2012, discussed later), allows the user to create sandboxes or customizable workspaces to support individual analysis routines.

A spectrum of interactions is available between the two extremes. Interaction techniques include changing visual scale (zooming), filtering, grouping and summarizing, and fetching details on demand.<sup>26</sup> Sarlin (2013, ch. 5) discusses many of the points between the extremes. These options include simple rendering refinements, data browsing, and exploratory data analysis, feature extraction, knowledge discovery in databases, and data mining. For macroprudential modeling, the functions should also include a range of domain-specific analytical and econometric techniques, such as those surveyed by Sarlin (2013, ch. 3) and Bisias, et al. (2012).


---

<sup>25</sup> Of course, fixed images are not analysis-free: They will inevitably be a snapshot or endpoint from some analytical process. The key requirement, for common knowledge communication and accountability, is that images be fixed. This constraint immediately makes other considerations much more salient, such as information density and graphic design, because interactive data browsing, summarization, details on demand, etc., are unavailable to the user.

<sup>26</sup> Pike, et al. (2009) refer to interaction as a cognitive act enabled by the computational tool but note this cognitive act does not take place exclusively through interaction with any single tool. For this reason, they argue the science of interaction encompasses a much broader concept than just the principles for creating interface widgets.

## Design Considerations for Visualizations of Financial Stability

Information visualization is not new to the domain of systemic supervision. However, such information visualizations cannot scale to large or high-dimensional data sets. The scale of the problem has been changing gradually for many years. For example, transaction-level data sets are beginning to become available to analysts for many markets. Each year, equity markets generate a few million daily closing prices for the universe of U.S. stocks. As high-frequency trading expands, trade-level data for the U.S. consolidated tape is producing roughly 10 times that many observations every day (Jones, 2013). These new, much larger data sets can be filtered or aggregated to make them accessible to legacy tools but to do so would mean discarding large amounts of potentially useful information. It is inevitable that new visualization tools will emerge to help address the larger data sets (Keim, Qu, and Ma, 2013; Fox and Hendler, 2011; Choo and Park, 2013).



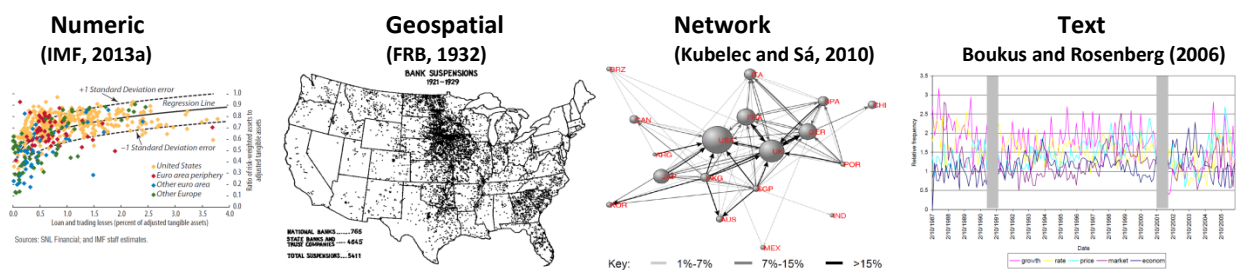
**Figure 11: How to Fail at Visualization**

The converse situation also poses a potential challenge. General-purpose new visualization libraries are now freely available (Bostock, et al., 2011), vastly reducing the cost of generating high-quality images and interactive displays. These new libraries also sharply reduce the cost of tossing off low-quality eye candy that attracts and engages the user but without effectively conveying any useful information, or even misleading the interpreter. The typical presumption when applying a generic display to existing data is that inputs are already well constructed and semantically useful — and the analyst should be interested in the image produced.

Unfortunately, many stability measures are new and poorly understood, and many of the new datasets are similarly unfamiliar and messy. Figure 11 shows how easy it is to distract — or

worse, actively mislead — users by incautiously tossing arbitrary data into a visualization tool. The left panel shows the importance of understanding the rules of perception (Ware, 2013). The right panel shows a mash-up of Consumer Price Index and gross domestic product in levels from the Federal Reserve Bank of St. Louis’s Federal Reserve Economic Data (FRED) tool (FRB-STL, 2013). FRED makes available thousands of clean, well-documented economic data series, together with a generic time-series charting tool. Unfortunately, the tool alone does not turn the user into a time-series econometrician; basic mistakes are simply easier to achieve quickly.

Sarlin (2013a, ch. 3) does a good job of surveying the current crop of financial-stability analytics in the context of a variety of visualization approaches, following the European Central Bank (ECB, 2010a) in decomposing financial stability risks according to origin (systematic vs. idiosyncratic), effect (simultaneous vs. sequential), and trigger (exogenous vs. endogenous), reducing eight categories to three broad forms: 1) endogenous build-up and unraveling of widespread imbalances; 2) exogenous aggregate shocks; and 3) contagion and spillover (Sarlin, 2013a, pp. 39–40). We focus here on a few highlights from that literature.

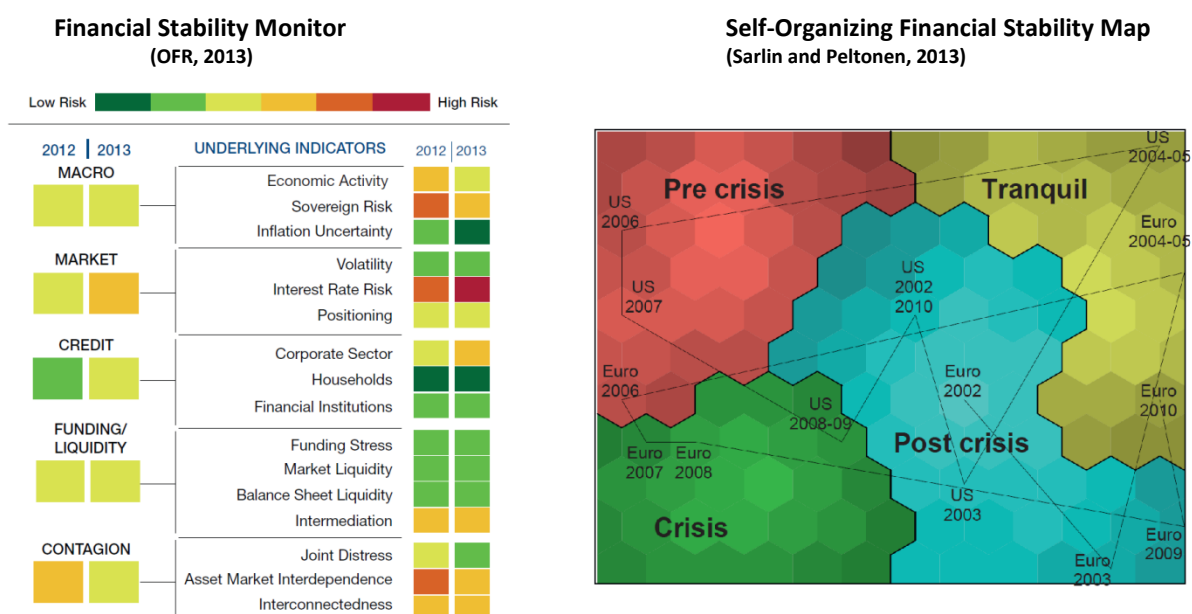


**Figure 12: Four Types of Financial Stability Data**

Given the complexity of the problem, the new data sets that are becoming available, and the growing range of systemic risk models (see Bisias, et al., 2012), some researchers and macroprudential supervisors have begun to use and advocate for visualizations to gain understanding.<sup>27</sup> The use of new visualizations often goes hand-in-hand with proposals to adopt new financial systemic risk measures or analytic techniques, such as network analysis. Zhang, et al. (2012, pp. 176-177), for example, identify four broad categories of data, each of which has special visualization considerations: 1) numeric; 2) geo-related; 3) network; and 4) text/web. The

<sup>27</sup> Instances are too numerous to list exhaustively. Some examples include Sarlin (2013a, 2013b), Sarlin and Peltonen (2013), Zhang, et al. (2012), Billio, et al. (2010), Flood, et al. (2012), Balakrishnan, et al. (2010), and Markose (2013).

data type of the dominant dimensions for comparison primarily distinguishes the four types. Figure 12 shows that all four types are relevant in financial stability analysis.

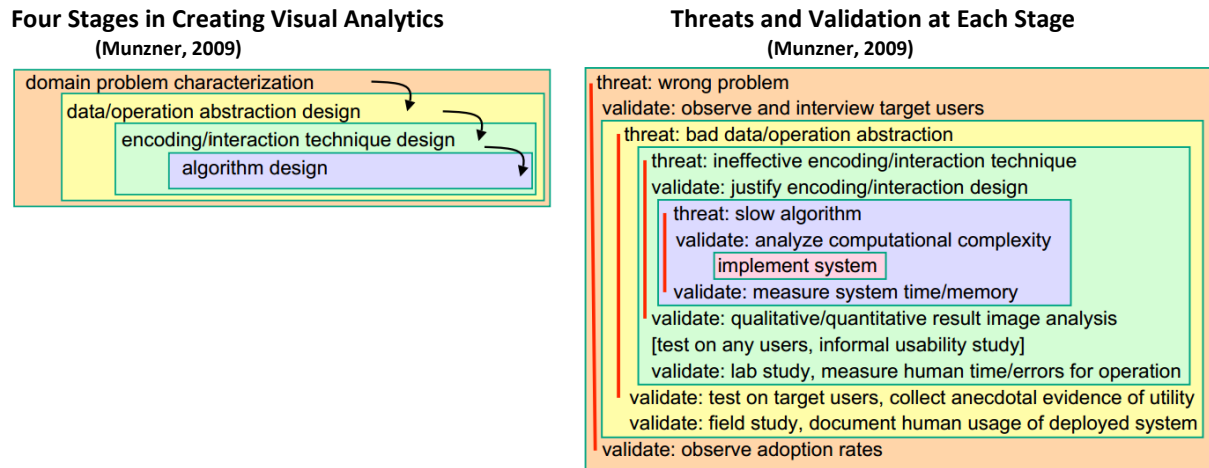


**Figure 13: Two Financial Stability Maps**

This four-part classification scheme is based on the raw data, the lowest level of the Data-Information-Knowledge-Wisdom hierarchy. What happens as model-based transformations move the data closer to “knowledge” is also important. The two approaches in Figure 13 illustrate recent attempts to represent systemic risk with numerous indicators simultaneously. Presenting multiple indicators jointly is sensible, because the financial system is multifaceted, so a higher dimensional data index is needed to capture everything. The left panel, for example, presents a financial stability heat map with five aggregate measures, each comprising several submeasures of financial systemic risk, (OFR, 2013, p. 11).

This visualization works well enough for high-level transparency, but like the right-hand panel of Figure 13, it would be less appropriate for an intensive sensemaking application. An analyst would likely wish to see the values of the individual submeasures in a sensemaking application. In a well-designed application, those values would likely be readily available, for example, via a double-click gesture to retrieve details on demand. Any attempt to show the disaggregated

submeasures on the present chart would likely make the map illegible.<sup>28</sup> Conversely, the self-organizing financial stability map in the right-hand panel of Figure 13, like the left-hand panel of Figure 6, is abstract and very information-dense.<sup>29</sup> As with IDL in Figure 6, a certain amount of specialized training is needed to achieve facility with manipulating the self-organizing map.



**Figure 14: Munzner's Nested Model for Visualization Design**

## Examples of Visual Analytics in Financial Stability Monitoring

In contrast to the examples in Section 2, the following examples focus on interactive visualization. The implementation process for augmenting visualizations with interactivity brings its own challenges. Munzner (2009) captures the lessons of experience in the nested model depicted in Figure 14. The basic design stages appear in the left panel. Note that the two outer stages, domain characterization and abstraction design, apply to visualization generally, Interaction comes into play significantly only at the third and fourth stages. A key point of the nesting, shown in the right panel, is that much of the validation needed at each stage to confirm threats have been adequately addressed cannot occur until nested implementations are complete and data are flowing through the system.

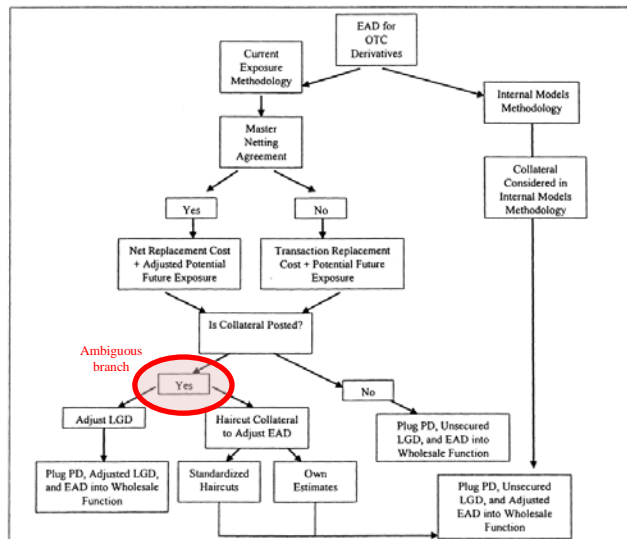
<sup>28</sup> Of course, precomputed detail maps might provide certain forms of supplemental drill-down in a similar format, but these would not be interactive. A visual analytics solution, in contrast, might allow the user to choose in real time among the various unions and intersections of several different drill-down dimensions, such as geographic region, industrial sector, etc.

<sup>29</sup> The Self-Organized Financial Stability Map uses a specialized artificial neural network, clustering country-level data on a wide range of indicators into groups interpreted as stages of the financial stability cycle. The approach uses a neighborhood function to preserve the topological properties of the high-dimensional input space. Sarlin (2013b) and Sarlin and Peltonen (2013) cover the details of the derivation of the self-organized map.



### Ad Hoc Flowchart (OCC-FRB-FDIC-OTS, 2006)

Figure 3—EAD and LGD for OTC Derivative Contracts



### Business Process Model and Notation

EAD and LGD for OTC Derivative Contracts

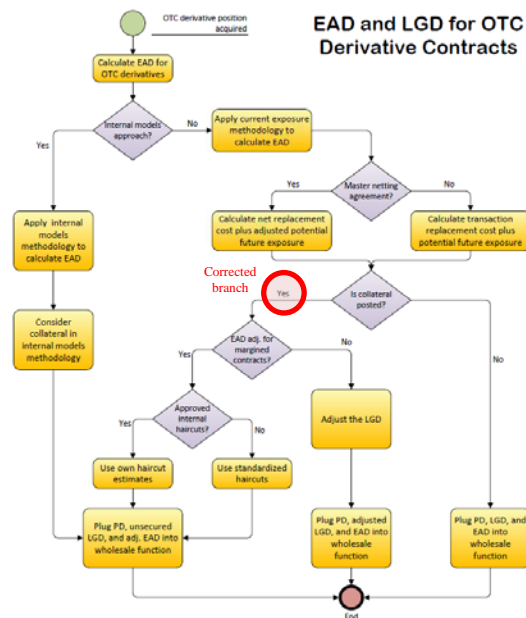


Figure 15: Formal Structure in a Process Diagram

The benefits of formal structure are apparent in Figure 15, which compares two versions of a proposed process for managing credit risk calculations under the Basel capital accord. The left panel comes from a formal decision-making process, the actual Notice of Proposed Rulemaking for a portion of the Basel II rule (OCC-FRB-FDIC-OTS, 2006, p. 55872). The right panel presents the identical workflow under the formalism of Business Process Model and Notation (often called BPMN; see OMG, 2011). The business process model provides a well-defined visual grammar for describing business processes. For example, compared to the relatively immature flowchart in the left panel, the business process model visually distinguishes process activities from decision gateways and their input and output flows. The discipline of the business process model avoids incoherencies such as the ambiguous “yes” box (left panel, circled in red), which diverges without explanation into two mutually exclusive pathways. The sequential activities depicted in rectangles in the right panel are atomic tasks with descriptions that adhere to the verb-plus-object structure. This rigid declarative syntax can facilitate automated implementation. The process model’s formalism connects directly to the Web Services Business Process Execution Language (or BPEL; see OASIS, 2007). This combination allows for a tight synchronization of the BPMN visualization with an automated behavioral implementation described by BPEL. This integration of user-friendly visualizations with machine-friendly

programming can reduce an important source of operational risk by facilitating straight-through processing. Users understanding the process described by the business process model flowchart can trust that the underlying implementation accurately follows the visual representation.

**Clockwise from Top Left: Heat Map, Search-by-Example, Keyword Graph, and Strings-and-Beads**  
(Chang, et al., 2007)

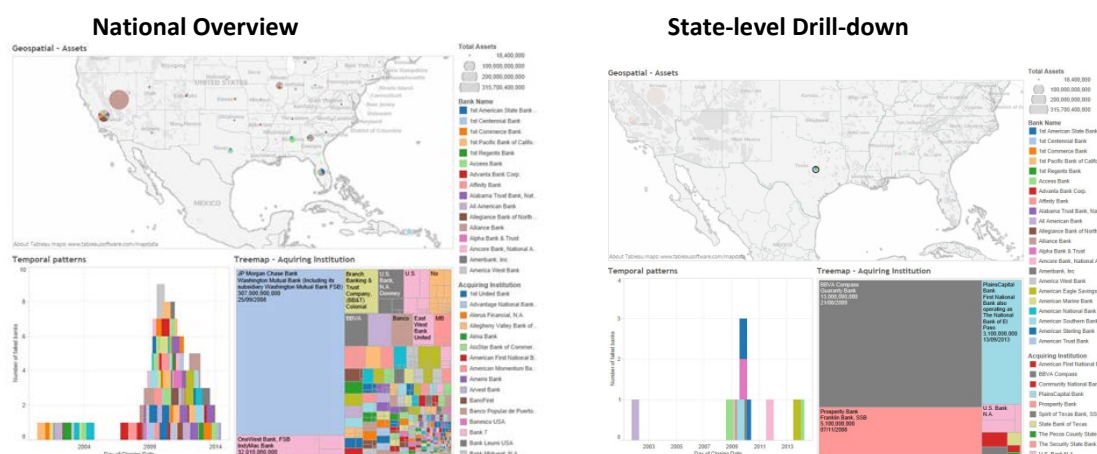


**Figure 16: Dashboard View of Wire-transfer Activity for a Bank**

Given the complexity of systemic analysis, it is often helpful to integrate various individual visualizations into dashboards that display multidimensional data in multiple coordinated views. Dashboards are mash-ups of diverse, coordinated perspectives on a collection of information, often optimized real-time monitoring (Few, 2006; MacNeil and Elmqvist, 2013). Figure 16, for example, shows a dashboard view of wire transactions, WireVis (Chang, et al., 2007). It allows interactive exploration of large numbers of categorical, time-varying, wire transaction data, as opposed to looking at each display in isolation. This dashboard displays a keyword network view, a heat map, a search-by-example tool, and a “strings-and-beads” display. The combination of the four visualizations emphasizes the relationships among accounts, time, and keywords within the transactions and offers the user not only a global overview of the data, but also the ability to compare, aggregate, and organize groups of transactions. This combination provides an efficient means of exploring, investigating, and analyzing the data, as well as drilling down

into and comparing individual records to make sense of the data and situations so as to enable informed decision making.

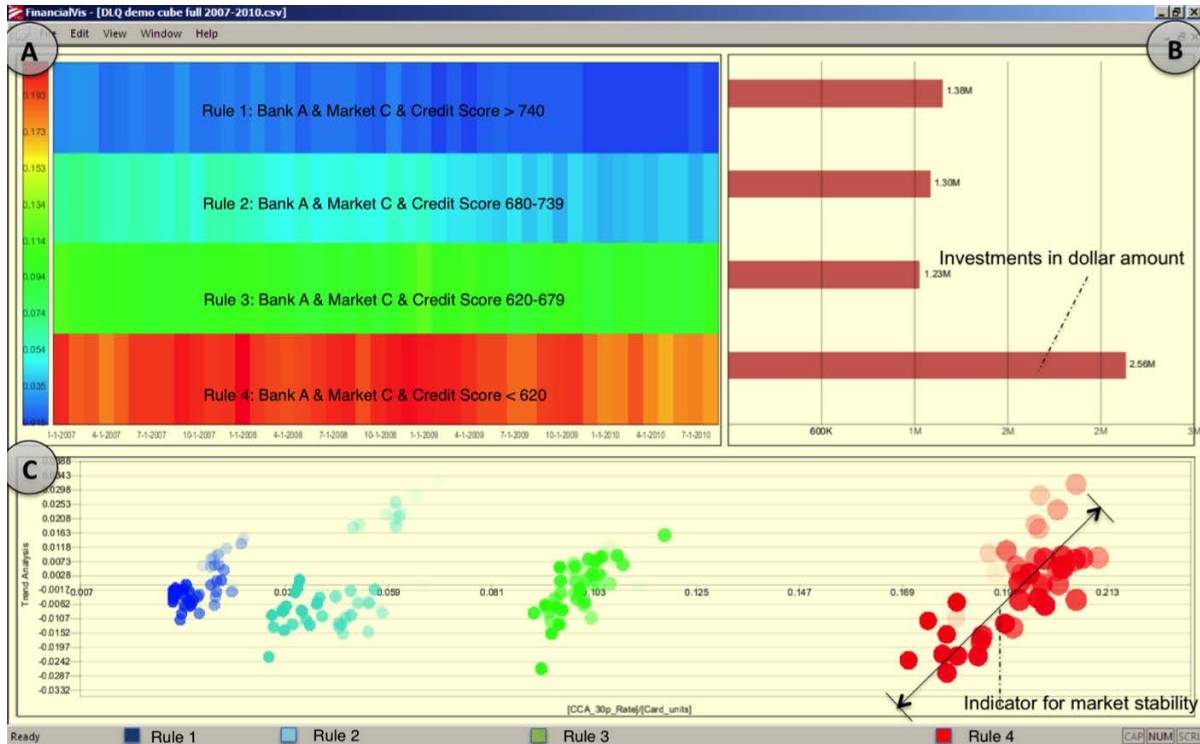
The largest bank failure in U.S. history occurred when Washington Mutual collapsed in 2008 with \$307 billion in assets. Loan defaults and other credit losses are dominant sources of risk for most banks and financial institutions, and credit assessment is crucial in maintaining a bank's financial stability. As Figure 2 shows, bank failure is a binary variable, readily comparable across institutions, and structural patterns typically emerge in episodes where failures are widespread. More recently, there were 519 failed banks from 2000 to 2014; some of the failed banks were acquired by other institutions.



**Figure 17: Interactive View of Failed Banks in the United States, 2000 to 2014 (FDIC data)**

Figure 17 (left panel) displays the 519 failed banks and their geographical locations. In this case, the dashboard is oriented toward sensemaking exercises by research analysts. The pie charts placed geographically represent the banks by their state locations, the colors represent the banks in question, and the size corresponds to the total assets. In contrast to Figure 2, which depicts bank failures before the advent of interstate branching, Figure 17 reveals that headquarters location has become more arbitrary in the modern era. The biggest pie chart corresponds to Nevada, which had 12 failed banks with combined total assets of more than \$315 billion, \$307 billion of which was from one large institution, Washington Mutual. The bottom left panel shows temporal information on when the banks failed while the treemap at the lower right groups failed banks by the institutions that acquired them. Acquirors are shown in different colors, with size representing total assets. The dashboard summarizes the location, size, and timing of bank failures. We can see that there were no bank failures during 2005 and 2006 and that some

institutions acquired more banks than others and with varying amounts of total assets. For example, JP Morgan Chase acquired Washington Mutual. 2010 had the most bank failures (157 banks with assets totaling \$96 billion).



**Figure 18: RiskVA**

As in Figure 2, the user can see at a glance that the number of failed banks varies across the states. The significant difference is the opportunity for interaction to explore the data. For example, clicking on any part of the dashboard highlights the corresponding features in the other parts of the dashboard through the application of brushing and linking techniques. In the right hand dashboard in Figure 17, clicking on Texas shows it had 12 failed banks, with the lower-right panel revealing all 12, including failure dates and the 10 institutions that acquired them, as well as the total assets of each of the banks and the overall value. The temporal history is shown on the lower left panel. Interaction allows comparisons of interest to the user, a crucial component of data exploration and hypothesis generation.

RiskVA (Figure 18) is a visual analytics system developed by University of North Carolina at Charlotte and Bank of America for consumer credit risk analysis. The system allows interactive data exploration and correlation and visually assists credit analysts by displaying coordinated

views of market fluctuations and temporal trends of the targeted credit product. Figure 18 is an overview of RiskVA, where panel (A) is the entity heat map, (B) shows the product comparison, and (C) is the trend analysis view. This display provides an interactive means for an analyst to examine the 30-day delinquency rate in a market across a range of consumers. Once again, the point is to encourage the exploration of financial information for knowledge and novel insight.

## 5. Conclusion and Directions for Future Research

Addressing the information processing challenges that have contributed to the global financial crisis remains a significant and unresolved challenge. Visual analytics offers a promising approach. Visual analytics combines the strength of analytical reasoning, the deep contextual knowledge of domain experts, and the unsurpassed pattern recognition capacity of the human visual system using interactive visual interfaces. Good visualizations work by presenting important facts, measured on relevant scales, and laid out so juxtaposition encourages comparison and reconciliation. Visual analytics augments this list with analytic algorithms and interaction tools that allow the user to steer the depiction, for example, applying Shneiderman's (1996) mantra, "overview first, zoom and filter, then details on demand."

The general principles of visualization and visual analytics are well understood, but there are always research opportunities when these tools are brought to bear in a new application domain. The challenge looms especially large for financial stability monitoring, where the size, scope, and changeability of the system is compounded by the abstract nature of the tasks and formative understanding of useful analytical approaches. In this context, we offer some preliminary and incomplete suggestions for research:

- ***Definition of core abstractions*** — Visualization tends to work best when the data share one or more measurable dimensions that form a basis for comparison. Ideally, the dimensions (definitions, not the specific values) will be invariant across observations. The ultimate goal is a mapping from the data and vocabulary of financial stability to the more generic domain of visual analytics. Bridging the semantic gap between domain abstractions and visualizations requires a deep understanding of both financial stability

analytics and a respect for implementation methodologies like the abstraction-decomposition hierarchy.

- ***Definition of canonical algorithms*** — Interactive visualization frequently calls for the calculation of derived attributes “on the fly” to supply characteristics for user-defined perspectives such as customized clusterings or filterings. For example, an analyst might filter out particular subsets of banks according to idiosyncratic values and then calculate the average liquidity coverage ratio or risk-weighted capital ratio for each subset. Similar to core data abstractions, there is a need for precise and semantically relevant algorithms, ideally with fast implementations, for embedding in visual analytics (Bisias, et al., 2012, offer a start by providing skeleton source code to accompany the models they describe in their survey).
- ***Publication of test data*** — Visualization researchers need data sets with relevant scope and content to prototype and test their tools. “Live” data are preferred to synthetic substitutes, although licensing and confidentiality concerns may predominate. Whether live or synthetic, sharing data implicitly raises issues of standardization, formatting and licensing.
- ***Development of evaluation techniques*** — Visualization tools can support a wide range of applications, including broad categories such as sensemaking or decision support, and more targeted purposes, such as representing semantic relationships in knowledge bases, depicting degrees of risk and uncertainty in financial data, or identifying gaps and quality issues in raw source data. As tooling emerges to address these various concerns, evaluation techniques will be needed to assess effectiveness. For example, decision-support tools for financial stability analysis will interact with users’ risk-based reasoning processes (see Oaksford, et al., 2012), but little is understood about how to assess whether that interaction is effective.

We have suggested some visualization approaches to financial stability monitoring and analysis, but research is needed to validate them. As a foundation for effectively applying visual analytics to meet these challenges, we underscore the need for stable abstractions and models related to financial stability and systemic risk and for approaches to address the continuing data availability and quality challenges.

In conclusion, we have discussed some of the possibilities as well as some of the pitfalls in applying visual analytics to the challenges of systemic financial stability monitoring. Much research remains to be done to further articulate the high- and low-level domain tasks, complete data and task abstractions, such as abstraction-decomposition hierarchies, develop visualizations and analyses to complement visualizations we have suggested could be incorporated into a visual analytics solution, and to enhance interaction techniques to aid exploration and analytic tasks. Validation and evaluation will also be needed. With further research and development, though not a panacea, visual analytics holds promise as a useful approach that could help financial systemic risk analysts meet the significant challenges associated with detecting, identifying, monitoring, and managing threats to global financial stability.

## 6. References

- Acharya, V., Pedersen, L., Philippon, T. and Richardson, M. (2010), “Measuring Systemic Risk,” Working paper, New York University.
- Achon, J. and Jamieson, G. (2003), “Work Domain Analysis of a Financial System: An Abstraction Hierarchy for Portfolio Management,” in: *Proceeding of the 22nd European Annual Conference on Human Decision Making and Control*, 103–109.
- Amengual, D. and Xiu, D. (2013), “Resolution of Policy Uncertainty and Sudden Declines in Volatility,” Working paper (13-78), Chicago Booth School of Business.
- Arias-Hernández R, Kaastra LT, Green TM, Fisher B (2011), “Pair analytics: capturing reasoning processes in collaborative VA.” In: R. Sprague (ed.), *Proceedings of the 44th Annual Hawaii International Conference on System Sciences*, 4-7 January 2011, Koloa, Kauai, Hawaii, IEEE Computer Society Press.
- Balakrishnan, S., V. Chu, M. Hernández, H. Ho, R. Krishnamurthy, S.X. Liu, J. Pieper, J. Pierce, L. Popa, C. Robson, L. Shi, I. Stanoi, E. Ting, S. Vaithyanathan, H. Yang (2010), “Midas: Integrating Public Financial Data,” in: *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data* (SIGMOD ’10), 1187–1190.
- Bank of England (BoE) (2013), “Financial Stability Report,” technical report, Bank of England. <http://www.bankofengland.co.uk/publications/Pages/fsr/2013/fsr34.aspx>
- Basel Committee on Banking Supervision (2011), “Basel III: A global regulatory framework for more resilient banks and banking systems,” technical report 189, Bank for International Settlements. <http://www.bis.org/publ/bcbs189.htm>
- Bennett, K. B., and Flach, J. M. (2011). *Display and Interface Design: Subtle Science, Exact Art*. Boca Raton: CRC Press, Taylor and Francis Group.
- Berner, R. (2011), “Testimony of Richard Berner, U.S. Department of the Treasury House Financial Services Subcommittee on Oversight and Investigations Hearing on ‘Oversight of the Office of Financial Research and the Financial Stability Oversight Council’,” Office of Financial Research, July 14, 2011. <http://financialservices.house.gov/uploadedfiles/071411berner.pdf>
- Billio, M., Getmansky, M., Lo, A. W. and Pelizzon, L. (2010), “Econometric measures of systemic risk in the finance and insurance sectors,” NBER working paper 16223, National Bureau of Economic Research.
- Bisias, D., Flood, M., Lo, A. and Valavanis, S. (2012), “A Survey of Systemic Risk Analytics,” *Annual Review of Financial Economics*, 4, 255–296.



- Bjørke, J. T., Nilsen, S. and Varga, M. J., (2010a), “Visualization of network structure by the application of hypernodes,” *International Journal of Approximate Reasoning*, 51, 275–293.
- \_\_\_\_\_, (2010b), “Visual analytics of data in Worldwide Incidents Tracking System and Alliance Data,” *NATO Workshop on Visualising Networks: Coping with Change and Uncertainty*, IST-093/RWS-015, Rome, New York, 19th – 21st October.
- Bjørke, J. T. and Varga, M. J., (2013), “Hypernode Annex F,” NATO STO Technical Report STO-TR-IST-085/RTG41 Interactive Visualisation of Network Dynamics.
- Blanchard, B. S., and W. J. Fabrycky (2010), *Systems Engineering and Analysis*, 5th Ed., Prentice Hall.
- Bloom, N. (2009), “The Impact of Uncertainty Shocks,” *Econometrica*, 77(3), 623–685.
- Board of Governors of the Federal Reserve System (FRB) (2013a), “Federal Open Market Committee,” Internet resource. <http://www.federalreserve.gov/monetarypolicy/fomc.htm>
- \_\_\_\_\_, (2013b), “Comprehensive Capital Analysis and Review: Assessment Framework and Results,” technical report. <http://www.federalreserve.gov/bankinfo/ccar-2013-results-20130314.pdf>
- \_\_\_\_\_, (2012a), “99th Annual Report, 2012,” technical report. <http://www.federalreserve.gov/publications/annual-report/files/2012-annual-report.pdf>
- \_\_\_\_\_, (2012b), “General Instructions for the Preparation of the Annual Capital Assessments and Stress Testing Information Collection (FR Y-14A),” Internet resource. [http://www.federalreserve.gov/reportforms/forms/FR\\_Y-14A20120930\\_f.zip](http://www.federalreserve.gov/reportforms/forms/FR_Y-14A20120930_f.zip)
- \_\_\_\_\_, (2012c), “General Instructions for the Preparation of the Annual Capital Assessments and Stress Testing Information Collection (FR Y-14Q/M),” Internet resource. [http://www.federalreserve.gov/reportforms/forms/FR\\_Y-14M20120930\\_f.zip](http://www.federalreserve.gov/reportforms/forms/FR_Y-14M20120930_f.zip)
- \_\_\_\_\_, (2007), “Monetary Policy Alternatives: Prepared for the Federal Open Market Committee (Bluebook),” technical report, FOMC. <http://www.federalreserve.gov/monetarypolicy/files/FOMC20070918bluebook20070913.pdf>
- \_\_\_\_\_, (1932), *Bank Suspensions in the United States, 1892-1931*, Federal Reserve Committee on Branch, Group, and Chain Banking.
- Bostock, M., Ogievetsky, V. and Heer, J. (2011), “D3: Data-Driven Documents,” *IEEE Transactions on Visualization and Computer Graphics*.
- Boukus, E. and Rosenberg, J. (2006), “The information content of FOMC minutes,” technical report, Federal Reserve Bank of New York. [http://ftp.ny.frb.org/research/economists/rosenberg/Boukus\\_and\\_Rosenberg\\_072006.pdf](http://ftp.ny.frb.org/research/economists/rosenberg/Boukus_and_Rosenberg_072006.pdf)

- Brammertz, W. (2013), The Office of Financial Research and Operational Risk, in: Victoria Lemieux, ed., *Financial Analysis and Risk Management: Data Governance, Analytics and Life Cycle Management*, Springer, 47–71.
- Brammertz, W., Akkizidis, I., Breymann, W., Entin, R. and Rustmann, M. (2009), *Unified Financial Analysis: the missing links of finance*, Wiley and Sons.
- Braswell, J. and Mark, R. (2013), “Banking and financial activities in the real economy,” in: M. Brose, M. Flood, D. Krishna and W. Nichols, eds., *Handbook of Financial Data and Risk Information, Volume I: Context*, Cambridge University Press.
- Burns, C. M. and Hajdukiewicz, J. R. (2004), *Ecological Interface Design*, CRC Press, Inc..
- Caballero, R. J. (2009), “The ‘Other’ Imbalance and the Financial Crisis,” Working paper No. 09-32, Massachusetts Institute of Technology.
- Caruana, J. (2012), “Interconnectedness and the importance of international data-sharing,” Speech by Mr Jaime Caruana, General Manager of the BIS, at the 3rd Swiss National Bank - International Monetary Fund conference on the international monetary system , Zurich, 8 May 2012, Bank for International Settlements. <http://www.bis.org/speeches/sp120730.htm>
- Chan, K. K. and Milne, A. (2013), “The Global Legal Entity Identifier System: Will it Deliver?” working paper, Loughborough University, August. [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2325889](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2325889).
- Chang, R., Ghoniem, M., Kosara, R., Ribarsky, W., Yang, J., Suma, E., Ziemkiewicz, C., Kern, D. and Sudjianto, A. (2007), “WireVis: Visualization of categorical, time-varying data from financial transactions,” in: *IEEE Symposium on Visual Analytics Science and Technology*, 2007, 155–162.
- Cherny, L. (2013), “A Fast-and-Dirty Intro to NetworkX (and D3),” technical report, Arnicas. <http://www.slideshare.net/arnicas/a-quick-and-dirty-intro-to-networkx-and-d3>
- Choo, J. and Park, H. (2013), “Customizing Computational Methods for Visual Analytics with Big Data,” *IEEE Computer Graphics and Applications*, 33(4), 22–28.
- Clark, D. R. (2010), “Understanding and Performance,” Internet resource, Big Dog & Little Dog’s Performance Juxtaposition. <http://www.nwlink.com/~donclark/performance/understanding.html>
- Commodity Futures Trading Commission, and Securities and Exchange Commission (2010), “Findings Regarding the Market Events of May 6, 2010: Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues,” technical report, CFTC/SEC. <http://www.sec.gov/news/studies/2010/marketevents-report.pdf>

- Chuang, J., Manning, C. D., and Heer, J. (2012) “‘Without the clutter of unimportant words’: Descriptive keyphrases for text visualization,” *ACM Transactions Computer-Human Interaction*, 19, 3, Article 19, October 2012.
- Crandall B, Klein G, Hoffman RR (2006), *Working minds: a practitioner’s guide to cognitive task analysis*, Bradford.
- Dervin, Brenda (1999), “Chaos, Order and Sense-making: A Proposed Theory for Information Design,” Ch. 3 in: R. Jacobson (ed.), *Information Design*, MIT Press, 35–57.
- Dick-Nielsen, J. (2009), “Liquidity Biases in TRACE,” *Journal of Fixed Income*, 19(2), 43–55.
- Durden, T. (2010), “Visualizing the Past of the Treasury Yield Curve, and Deconstructing the Great Confusion Surrounding Its Future,” Internet resource, ZeroHedge.  
<http://www.zerohedge.com/article/visualizing-past-treasury-yield-curve-and-deconstructing-great-confusion-surrounding-its-fut>
- Engel, P., Hamscher, W., Shuetrim, G., von Kannon, D. and Wallis, H. (2013), “Extensible Business Reporting Language (XBRL) 2.1: Recommendation 31 December 2003 with errata corrections to 20 February 2013,” technical report, XBRL International.  
<http://www.xbrl.org/Specification/XBRL-2.1/REC-2003-12-31/XBRL-2.1-REC-2003-12-31+corrected-errata-2013-02-20.html>
- European Central Bank (ECB) (2010a), “Analytical models and tools for the identification and assessment of systemic risks,” *Financial Stability Review*, 138–146.
- \_\_\_\_\_ (2010b), “Financial networks and financial stability,” *Financial Stability Review*, 155–160.
- Federal Reserve Bank of St. Louis (FRB-STL) (2013), “Federal Reserve Economic Data (FRED),” Internet resource. <http://research.stlouisfed.org/fred2/>
- Few, S. (2006), *Information Dashboard Design: The Effective Visual Communication of Data*, O’Reilly.
- Financial Stability Board (2009), “Guidance to Assess the Systemic Importance of Financial Institutions, Markets and Instruments: Initial Considerations,” technical report, Financial Stability Board. [http://www.financialstabilityboard.org/publications/r\\_091107c.pdf](http://www.financialstabilityboard.org/publications/r_091107c.pdf)
- Financial Stability Board and International Monetary Fund (FSB-IMF) (2012), “The Financial Crisis and Information Gaps: Progress Report on the G-20 Data Gaps Initiative: Status, Action Plans, and Timetables,” technical report.  
<http://www.imf.org/external/np/g20/pdf/093012.pdf>
- \_\_\_\_\_ (2009), “The financial crisis and information gaps,” technical report.  
<http://www.imf.org/external/np/g20/pdf/102909.pdf>

- Financial Stability Oversight Council (2013), “FSOC 2013 Annual Report,” technical report, FSOC. [http://www.treasury.gov/initiatives/fsoc/Documents/FSOC 2013 Annual Report.pdf](http://www.treasury.gov/initiatives/fsoc/Documents/FSOC%2013%20Annual%20Report.pdf)
- Flood, M., Katz, J., Ong, S. and Smith, A. (2013), “Cryptography and the Economics of Supervisory Information: Balancing Transparency and Confidentiality,” OFR Working Paper #0011, Office of Financial Research. [http://www.treasury.gov/initiatives/ofr/research/Documents/OFRwp0011\\_FloodKatzOngSmith\\_CryptographyAndTheEconomicsOfSupervisoryInformation.pdf](http://www.treasury.gov/initiatives/ofr/research/Documents/OFRwp0011_FloodKatzOngSmith_CryptographyAndTheEconomicsOfSupervisoryInformation.pdf)
- Flood, M., Mendelowitz, A. and Nichols, W. (2012), Monitoring Financial Stability in a Complex World, in: V. Lemieux, ed., *Financial Analysis and Risk Management: Data Governance, Analytics and Life Cycle Management*, Springer Verlag.
- Flood, M. D., Kyle, A. and Raschid, L. (2010), “Knowledge Representation and Information Management for Financial Risk Management: Report of a Workshop: Arlington, VA, July 21-22, 2010,” technical report, U. of Maryland. <http://irix.umiacs.umd.edu/docs/FIWreport-FINAL.pdf>
- Fox, P. and Hendler, J. (2011), “Changing the Equation on Scientific Data Visualization,” *Science*, 331, 705–708.
- Frické, M. (2009), “The Knowledge Pyramid: A Critique of the DIKW Hierarchy,” *Journal of Information Science*, 35(2), 131–142.
- Group of Ten (G-10) (2001), “Report on Consolidation in the Financial Sector: Chapter III. Effects of consolidation on financial risk,” technical report, International Monetary Fund. <http://www.imf.org/external/np/g10/2001/01/eng/index.htm>
- Ginges, J., and Cairns, D. (2000). “Social representation of multiculturalism: A faceted analysis.” *Journal of Applied Social Psychology*, 30(7), 1345–1370.
- Gotz, D., and Zhou, M. (2008), “Characterizing users’ visual analytic activity for insight provenance,” in *IEEE Visual Analytics Science and Technology*, 123–130.
- Guttman, L. (1959). “Introduction to facet design and analysis.” *Proceedings of the 15th International Congress of Psychology in Brussels*, North-Holland, 130–132.
- Haynes, R., Paddrik, M., and Rajan, S. (2014), “Visual Network Analysis in the Regulation of Financial Systemic Risk, working paper, Office of Financial Research.
- Heer, J., and M. Agrawala (2008), “Design considerations for collaborative visual analytics,” *Information Visualization*, 7, 49–62.
- Heer, J., and Shneiderman, B. (2012). “Interactive Dynamics for Visual Analytics,” *Communications of the ACM*, 55(4), 45–54.

- Hey, J. (2004), “The Data, Information, Knowledge, Wisdom Chain: The Metaphorical link,” working paper,  
[http://web.archive.org/web/20071202033948/http://ioc.unesco.org/Oceanteacher/OceanTeacher2/02\\_InfTehSciCmm/DIKWchain.pdf](http://web.archive.org/web/20071202033948/http://ioc.unesco.org/Oceanteacher/OceanTeacher2/02_InfTehSciCmm/DIKWchain.pdf).
- Hunt, J. P., Stanton, R. and Wallace, N. (2013), “US residential-mortgage transfer systems: a data-management crisis,” in: M. Brose, M. Flood, D. Krishna and W. Nichols, eds., *The Handbook of Financial Data and Risk Information, Volume II: Software and Data*, Cambridge U. Press, pp. 85–132.
- Hutchins, E. (1995), *Cognition in the Wild*, MIT Press.
- International Monetary Fund (IMF) (2013a), “IMF Data Mapper,” Internet resource.  
<http://www.imf.org/external/datamapper/index.php>
- \_\_\_\_\_ (2013b), “Acute Risks Reduced: Actions Needed to Entrench Financial Stability,” *Global Financial Stability Report*, 1–56.
- \_\_\_\_\_ (2011), “Macroprudential Policy: An Organizing Framework,” technical report, March, <http://www.imf.org/external/np/pp/eng/2011/031411.pdf>.
- Jones, C. M. (2013), “What do we know about high-frequency trading?” technical report, Columbia Business School.
- Joyce, J. P. and Lapinsky, G. W. (1983), “A history and overview of the safety parameter display system concept,” *IEEE Transactions on Nuclear Science*, 30(1), 744–749.
- Kadivar, N., Chen, V., Dunsmuir, D., Lee, E., Qian, C., Shaw, C., and Woodbury, R. (2009), “Capturing and supporting the analytic process,” *IEEE Visual Analytics Science and Technology*, 131–138.
- Kahneman, D. (2011), *Thinking Fast and Slow*, Farrar, Straus and Giroux.
- Kapadia, S., Drehmann, M., Elliott, J. and Sterne, G. (2009), “Liquidity Risk, Cash Flow Constraints, and Systemic Feedbacks,” technical report, Bank of England.
- Keim, D. A., Kohlhammer, J., Ellis, G. and Mansmann, eds. (2011), *Mastering the Information Age – Solving Problems with Visual Analytics*, Florian Mansmann.
- Keim, D., Mansmann, F., Schneidewind, J. and Ziegler, H. (2006), Challenges in Visual Data Analysis, in: *Tenth International Conference on Information Visualization*, 2006, pp. 9–16.
- Keim, D., Qu, H. and Ma, K.-L. (2013), “Big Data Visualization,” *IEEE Computer Graphics and Applications*, 33(4), 20–21.

- Keys, B. J., Mukherjee, T., Seru, A. and Vig, V. (2010), “Did securitization lead to lax screening? Evidence from subprime loans,” *Quarterly Journal of Economics*, 125(1), 307–362.
- Khandani, A. E., Kim, A. J. and Lo, A. W. (2010), “Consumer credit-risk models via machine-learning algorithms,” *Journal of Banking and Finance*, 34(11), 2767–2787.
- Klein, G., Moon, B. M., and Hoffman, R. R. (2006a), “Making Sense of Sensemaking 1: Alternative Perspectives,” *IEEE Intelligent Systems*, 21(4), July/August, 70–73.
- \_\_\_\_\_ (2006b), “Making Sense of Sensemaking 2: A Macrocognitive Model,” *IEEE Intelligent Systems*, 21(5), September/October, 88–92.
- Koning, J. P. (2012), “Managed Currency: The People’s Bank of China balance sheet since 2002,” Internet resource, Financial Graph and Art. <http://jpkoning.blogspot.ca/2012/11/data-visualization-peoples-bank-of.html>
- Kubelec, C. and Sá, F. (2010), “The geographical composition of national external balance sheets: 1980-2005,” Bank of England Working Paper 384. <http://tna.europarchive.org/20101014143014/http://www.bankofengland.co.uk/publications/workingpapers/wp384.pdf>
- Lemieux, V. L. (2013), “Records and Information Management for Financial Analysis and Risk Management: An Introduction,” Ch. 1 in: V. Lemieux, ed., *Financial Analysis and Risk Management: Data Governance, Analytics and Life Cycle Management*, Springer Verlag, 1–14.
- Lemieux, V.L. and Dang T., “Building Accountability for Decision Making into Cognitive Systems,” in A. Rocha, A. Correia, T. Wilson, and K. Stroetmann, eds., *Advances in Information Systems and Technologies*, Springer-Verlag, 2013, 575-586.
- Lemieux, V. L., Phillips P., Bajwa H. S., and Li C. (2014), “Applying Visual Analytics to the Global Legal Entity Identifier System to Enhance Financial Transparency,” conference presentation, Conference on Data Standards, Information, and Financial Stability, Loughborough University, April 11-14.
- Lemieux, V. L., Fisher, B. and Dang, T. (2014), “The visual analysis of financial data,” in: M. Brose, M. Flood, D. Krishna and W. Nichols, eds., *The Handbook of Financial Data and Risk Information: Volume 2: Software and Data*, Cambridge U. Press, 279–326.
- Lemieux, V.L., Rahmdel, P.S, Walker, R., Flood, M., and Wong, W., “Clustering Techniques and Their Effects on Portfolio Formation and Risk Analysis,” Data Science for Macro-Modeling Workshop, ACM SIGMOD, June 27, 2014 (forthcoming).



- MacNeil, S. and N. Elmqvist (2013), “Visualization Mosaics for Multivariate Visual Exploration,” *Computer Graphics Forum*, 14(3), 285–305.
- Markose, S. (2013), “Systemic risk analytics: A data-driven multi-agent financial network (MAFN) approach,” *Journal of Banking Regulation*, 14, 285–305.
- Memisevic, R., Sanderson, P., Choudhoury, S., and Wong, W. (2005), “Work domain analysis and ecological interface design for hydropower system monitoring and control,” *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, 4, (Hawaii, USA, 10-12 October 2005), 3580–3587.
- Mishkin, F. S. (2007), “Systemic Risk and the International Lender of Last Resort,” technical report, Board of Governors of the Federal Reserve, Speech delivered at the Tenth Annual International Banking Conference, Federal Reserve Bank of Chicago, September 28, 2007. <http://www.federalreserve.gov/newsevents/speech/mishkin20070928a.htm>
- Mohammed, S. (2001), “Toward an Understanding of Cognitive Consensus in a Group Decision-Making Context,” *Journal of Applied Behavioral Science*, 37(4), December , 408–425.
- Moussa, A. (2011), “Contagion and Systemic Risk in Financial Networks,” PhD thesis, Columbia University. <http://academiccommons.columbia.edu/catalog/ac:131474>
- Munzner, T. (2009), “A Nested Model for Visualization Design and Validation,” *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 15(6), 921-928. <http://www.cs.ubc.ca/labs/imager/tr/2009/NestedModel/>
- Nguyen, P., Kai Xu, Rick Walker and B. L. William Wong (2014a – submitted). “SchemaLine: Timeline Visualization for Sense Making.” Submitted to Information Visualisation.
- \_\_\_\_\_ (2014b – submitted). “TimeSets: A Visualisation with Set Relations.” Submitted to the IEEE Vis 2014.
- North Atlantic Treaty Organization (NATO) (2013). “NATO IST-117 Visualization for Analysis Workshop,” Shrivenham, U.K.
- Oaksford, L., Chater, N., and Stewart, N. (2012), “Reasoning and decision-making,” Ch. 7 in K. Frankish and W. M. Ramsey, eds., *The Cambridge Handbook of Cognitive Science*, Cambridge University Press, 131-150.
- Object Management Group (OMG) (2011), Business Process Model and Notation (BPMN), Version 2.0, OMG standard, January. <http://www.omg.org/spec/BPMN/2.0>
- Office of Financial Research (OFR) (2013), Annual Report, 2013. [http://www.treasury.gov/initiatives/ofr/about/Documents/OFR\\_AnnualReport2013\\_FINAL\\_12-17-2013.pdf](http://www.treasury.gov/initiatives/ofr/about/Documents/OFR_AnnualReport2013_FINAL_12-17-2013.pdf)

Office of the Comptroller of the Currency, Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation and Office of Thrift Supervision (OCC-FRB-FDIC-OTS) (2011), “Agency Information Collection Activities; Submission for OMB Review; Joint Comment Request,” *Federal Register*, 76(130), 39981–39989.

\_\_\_\_ (2006), “Risk-Based Capital Standards: Advanced Capital Adequacy Framework,” Notice of Proposed Rulemaking.  
[http://www.federalreserve.gov/GeneralInfo/Basel2/NPR\\_20060905/NPR/Basel\\_II\\_NPR.pdf](http://www.federalreserve.gov/GeneralInfo/Basel2/NPR_20060905/NPR/Basel_II_NPR.pdf)

Office of Financial Research (OFR) (2013), *2013 Annual Report*.  
<http://www.treasury.gov/initiatives/ofr/about/Pages/2013-Annual-Report.aspx>

OFR Financial Research Advisory Committee (OFR-FRAC) (2013), “Financial Services and Risk Management Subcommittee Recommendations, Part 2,” technical report, Office of Financial Research.  
[http://www.treasury.gov/initiatives/ofr/about/Documents/Financial%20Services%20and%20Risk%20Management%20Subcommittee\\_Recommendations\\_Part\\_2.pdf](http://www.treasury.gov/initiatives/ofr/about/Documents/Financial%20Services%20and%20Risk%20Management%20Subcommittee_Recommendations_Part_2.pdf)

Organization for the Advancement of Structured Information Standards (OASIS) (2007), Web Services Business Process Execution Language (BPEL), Version 2.0, OASIS standard, April.  
<http://docs.oasis-open.org/wsbpel/2.0/OS/wsbpel-v2.0-OS.pdf>

Pástor, L. and Veronesi, P. (2013), “Political Uncertainty and Risk Premia,” *Journal of Financial Economics*, 110(3), 520–545.

Pike, W. A., Stasko, J., Chang, R. and O’Connell, T. A. (2009), “The science of interaction,” *Information Visualization*, 8(4), 263–274.

Pileggi, H., Stolper, C. D., Boyle, J. M. and Stasko, J. T. (2012), “SnapShot: Visualization to Propel Ice Hockey Analytics,” *IEEE Transactions on Visualization and Computer Graphics*, 18(12), 2819–2828.

Quax, R., Kandhai, D. and Sloot, P. M. A. (2013), “Information dissipation as an early-warning signal for the Lehman Brothers collapse in financial time series,” *Scientific Reports*, 3, 1–7.

Rajaraman, Anand, Jure Leskovec, and Jeffrey D. Ullman (2014), Mining of Massive Datasets, manuscript, Stanford U. <http://infolab.stanford.edu/~ullman/mmds/book.pdf>

Rasmussen, J., Pejtersen, A. M. and Goodstein, L. P. (1994), *Cognitive systems engineering*, Wiley and Sons.

Reason, J. (1987), “The Chernobyl errors,” *Bulletin of the British Psychological Society*, 40, 201–206.



- Roberts, J. C., Keim, D., Hanratty, T., Rowlingson, R., Hall, M., Jacobson, Z., Lavigne, V., Rooney, C. and Varga, M. (2014), "From Ill-defined Problems to Informed Decisions." EuroVis Workshop on Visual Analytics, U.K.
- Rodrigues, E. M., N. Milic-Frayling, M. Smith, B. Shneiderman, and D. Hansen (2011), "Group-In-a-Box Layout for Multi-faceted Analysis of Communities," in: *Proceedings of the IEEE Third International Conference on Privacy, Security, Risk and Trust*, 354–361.
- Rosengren, E. S. (2010), "Asset Bubbles and Systemic Risk," technical report, Federal Reserve Bank of Boston, Speech delivered at the Global Interdependence Center's Conference on "Financial Interdependence in the World's Post-Crisis Capital Markets," Philadelphia, March 3, 2010. <http://www.bos.frb.org/news/speeches/rosengren/2010/030310/030310.pdf>
- Rowley, J. (2007), "The wisdom hierarchy: representations of the DIKW hierarchy," *Journal of Information Science*, 33(2), 163–180.
- Russell, D. M., Stefik, M. J., Pirolli, P., and Card, Stuart K. (1993), "The cost structure of sensemaking," *Proceedings of the INTERACT'93 and CHI'93 Conference on Human Factors in Computing Systems*, 269–276.
- Sanderson, P., Wong, W., Choudhury, S., and Memisevic, R. (2003), "Hydro Scheme Control in a Deregulated Environment: Cognitive Work Models and Design Implications," *Human Factors and Ergonomics Society (HFES), 47th Annual Meeting* (Denver, Colorado), 458–462.
- Sarkar, I. N. (2013), *Methods in Biomedical Informatics: A Pragmatic Approach*, Academic Press.
- Sarlin, P. (2013a), "Mapping Financial Stability," PhD thesis, Åbo Akademi University. [http://tuus.fi/publications/view/?pub\\_id=phdSarlin\\_Peter13a](http://tuus.fi/publications/view/?pub_id=phdSarlin_Peter13a)
- \_\_\_\_\_ (2013b), "Exploiting the self-organizing financial stability map," *Engineering Applications of Artificial Intelligence*, 26(5–6), May–June 2013, 1532–1539.
- \_\_\_\_\_ (2014), "Macroprudential oversight, risk communication and visualization," *working paper*, Goethe University, March.
- Sarlin, P. and Peltonen, T. A. (2013), "Mapping the State of Financial Stability," *Journal of International Financial Markets, Institutions and Money*, 26, 46–76.
- Savikhin, A. C. (2013), "The Application of Visual Analytics to Financial Decision-Making and Risk Management: Notes from Behavioral Economics," Ch. 5 in: V. Lemieux, ed., *Financial Analysis and Risk Management: Data Governance, Analytics and Life Cycle Management*, Springer Verlag, 101–114.
- Schich, M. (2010), "Revealing Matrices," in: J. Steele and N. Iliinsky, eds., *Beautiful Visualization: Looking at Data through the Eyes of Experts*, O'Reilly Press, 227–254.

- Schraagen JM, Chipman SF, Shalin VL (eds) (2000), *Cognitive task analysis*, Lawrence Erlbaum Associates.
- Schwabish, J. (2014), “An Economist’s Guide to Visualizing Data,” *Journal of Economic Perspectives*, 28(1), 209-234.
- Securities and Exchange Commission (SEC) (2013), “MIDAS Market Information Data Analytics System,” technical report. <http://www.sec.gov/marketstructure/midas.html>
- \_\_\_\_\_ (2012), “Consolidated Audit Trail,” technical report. <http://www.sec.gov/rules/final/2012/34-67457.pdf>
- Shneiderman, B. (1996), “The eyes have it: A task by data type taxonomy for information visualizations,” *Proceedings of the IEEE Symposium on Visual Languages*, 1996, 336–343. [http://drum.lib.umd.edu/bitstream/1903/5784/1/TR\\_96-66.pdf](http://drum.lib.umd.edu/bitstream/1903/5784/1/TR_96-66.pdf)
- Shneiderman, B., C. Plaisant, M. Cohen, and S. Jacobs (2009), *Designing the User Interface: Strategies for Effective Human-Computer Interaction*, 5th ed., Prentice Hall.
- Soramäki, K., Bech, M. L., Arnold, J., Glass, R. J. and Beyeler, W. E. (2007), “The topology of interbank payment flows,” *Physica A*, 379, 317–333.
- Sundström, G. A. and Hollnagel, E. (2011), *Governance and Control of Financial Systems: A Resilience Engineering Perspective*, Ashgate Publishing.
- Tatu, A., Zhang, L., Bertini, E., Schreck, T., Keim, D., Bremm, S. and von Landesberger, T. (2012), “ClustNails: Visual analysis of subspace clusters,” *Tsinghua Science and Technology*, 17(4), 419–428.
- Thomas, J. J. and Cook, K. A. (2005), *Illuminating the Path: The Research and Development Agenda for Visual Analytics*, National Visualization and Analytics Center.
- Trickett, S.B., J.G. Trafton, L. Saner, and C.D. Schunn (2007), “I don't know what's going on there: the use of spatial transformations to deal with and resolve uncertainty in complex visualizations.” In: Lovett MC, Shah P (eds.) *Thinking with data*, Lawrence Erlbaum Associates, 65-86.
- Treisman, A. (1985), “Preattentive Processing in Vision,” *Computer Vision, Graphics and Image Processing*, 31(2), 156–177.
- Tufte, Edward R. (2001), *The Visual Display of Quantitative Information*, 2nd edition, Graphics Press. [http://www.edwardtufte.com/tufte/books\\_vdqi](http://www.edwardtufte.com/tufte/books_vdqi)
- \_\_\_\_\_ (1990), *Envisioning Information*, Graphics Press. [http://www.edwardtufte.com/tufte/books\\_ei](http://www.edwardtufte.com/tufte/books_ei)

- van Wijk, J. J. (2005), "The value of visualization," *IEEE Visualization 2005*, 79–86.
- Vicente, K. J. (1999), *Cognitive work analysis: Toward safe, productive, and healthy computer-based work*, CRC Press.
- Vicente, K. J., and Rasmussen, J. (1992). "Ecological Interface Design: Theoretical Foundations." *IEEE Transactions on Systems, Man, and Cybernetics*, 22(4), 589–605.
- Vuckovic, A., Sanderson, P., Neal, A., Gaukrodger, S., and Wong, B. L. W. (2013), "Relative Position Vectors: An Alternative Approach to Conflict Detection in Air Traffic Control," *Human Factors*, published on-line before print 28 March 2013.
- Wang, X., Jeong, D., Chang, R. and Ribarsky, W. (2012), "RiskVA: A Visual Analytics System for Consumer Credit Risk Analysis," *Tsinghua Science and Technology*, 17(4), 440–451.
- Ware, C. (2013), *Information Visualization: Perception for Design*, Morgan Kaufmann.
- Watts, K. A. (2009), "Proposing a Place for Politics in Arbitrary and Capricious Review," *Yale Law Journal*, 119(1), 2–85.
- Weick, Karl E. (1995), *Sensemaking in Organizations*, Sage Publications.
- Wilkinson, Leland (2005), *The Grammar of Graphics, 2nd edition*, Springer Verlag.
- Walker, R., Slingsby, A., Dykes, J., Xu, K., Wood, J., Nguyen, P., et al. (2013). "An Extensible Framework for Provenance in Human Terrain Visual Analytics." *IEEE Transactions on Visualization and Computer Graphics*, 19(12), 2139-2148.
- Wong, B. W., Rozzi, S., Boccalatte, A., Gaukrodger, S., Amaldi, P., Fields, B., Loomes, M. and Martin, P. (2007), "3D-in-2D Displays for ATC," *6th EUROCONTROL Innovative Research Workshop*, 42–62.
- Woods, D. D., and Roth, E. M. (1988). Cognitive Systems Engineering. In M. Helander (Ed.), *Handbook of Human-Computer Interaction*: Elsevier Science Publishers B.V. (North-Holland), 3-43.
- Woods, D. D. and Hollnagel, E. (2006), *Joint cognitive systems: Patterns in cognitive systems engineering*, CRC Press.
- World Bank (2013a), "DataBank," Internet resource.  
<http://databank.worldbank.org/data/home.aspx>
- \_\_\_\_\_ (2013b), "Data Visualizer," Internet resource.  
<http://devdata.worldbank.org/DataVisualizer/>

- Yi, J. S., ah Kang, Y., Stasko, J. T. and Jacko, J. A. (2007), “Toward a deeper understanding of the role of interaction in information visualization,” *IEEE Transactions on Visualization and Computer Graphics*, 13(6), 1224–1231.
- Zhang, L., Stoffel, A., Behrisch, M., Mittelstadt, S., Schreck, T., Pompl, R., Weber, S., Last, H. and Keim, D. (2012), Visual analytics for the big data era—A comparative review of state-of-the-art commercial systems, in: *IEEE Conference on Visual Analytics Science and Technology (VAST)*, 173–182.